

Conservation Priorities and Environmental Offsets: Markets for Florida Wetlands[†]

By DANIEL ARONOFF AND WILL RAFEY*

We introduce an empirical framework for valuing markets in environmental offsets. Using newly collected data on wetland conservation and offsets, we apply this framework to evaluate a set of decentralized markets in Florida, where land developers purchase offsets from long-lived producers who restore wetlands over time. We find offsets led to substantial private gains from trade, creating \$2.4 billion of net surplus from 1995 to 2020 relative to direct conservation. Offset trading also generated new hydrological externalities. A locally differentiated Pigouvian tax would have prevented \$1.6 billion of new flood damage while preserving more than two-thirds of the private gains from trade. (JEL D47, D62, H23, Q24, Q54, Q57, R14)

Environmental offsets—contracts to remediate or restore the environment in lieu of direct abatement or conservation—play an increasingly central role in modern environmental regulation. Offset markets can create private gains from trade relative to more commonly used conservation mandates, but equilibrium outcomes in such markets will not be efficient unless regulators can account perfectly for the social value of offsets. In particular, while offsets can provide flexibility to conserve a public good at lower cost, they raise concerns when they cannot (or do not) substitute for all dimensions of the original public good.

This paper introduces an empirical framework for environmental market design in the presence of these two potentially competing concerns. A regulator specifies a conservation objective to preserve the existing stock of a public good. A set of potential producers access restoration opportunities that differ in cost as well as

*Aronoff: MIT Department of Economics and MIT Laboratory for Economic Analysis and Design (daronoff@mit.edu); Rafey: UCLA Department of Economics and NBER (rafey@econ.ucla.edu). Katja Seim was the coeditor for this article. We thank Mark Serrahn for invaluable research assistance and Jennifer Thomson and Alex Preisser for invaluable data access. We thank numerous seminar participants at various universities, as well as Steve Berry, Hannah Druckenmiller, and Stephen Ryan for serving as discussants for this paper at the 2024 NBER IO Winter Meetings, the 2023 NBER EEE Summer Institute, and the 2025 North American Winter Meeting of the Econometric Society. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program (Grant No. 1122374) and the George and Obie Schultz Fund. This work used computational and storage services associated with the Hoffman2 Cluster, which is operated by the UCLA Office of Advanced Research Computing's Research Technology Group. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or any of the other institutions listed. Our referees and coeditor provided many helpful comments. All errors are our own. Copyright Will Rafey and Daniel Aronoff, 2025. All rights reserved. None of this content can be used for model training or any other purpose by any software, computer program, and/or web application without the authors' express prior permission.

[†]Go to <https://doi.org/10.1257/aer.20231016> to visit the article page for additional materials and author disclosure statement(s).

location. Producers undertake long-run restoration activities, receive offset credits from the regulator, and sell offsets to entities seeking to deplete the public good. Offsets contribute to the regulator's conservation objective and may also have other environmental consequences. When estimated with data on offset producers and trade flows, the model allows us to recover the private gains from trade in offsets, measure the environmental outcomes from trade, and predict counterfactual gains from trade and environmental outcomes under alternative market designs.

We apply this framework to value a new set of decentralized markets for protected wetlands. Wetlands deliver a range of environmental benefits, including biodiversity, water purification, carbon sequestration, and flood protection.¹ At the same time, their preservation precludes competing land uses—such as housing, agriculture, or infrastructure—that may create private value. Federal and state environmental laws negotiate these trade-offs in the United States by mandating “No Net Loss” in existing wetlands. These rules allow development on wetlands if the loss is “offset” by an equal gain on other wetlands in the same region. This legal framework involves long-lived wetlands producers, who build or restore permanent wetlands on private land (“wetland mitigation banks”) to produce certified offsets, which they then sell to landowners developing protected wetlands.

To analyze these markets, we obtain new data on markets for wetland offsets in Florida, where 29 percent of land by area is wetlands and real estate comprised nearly one-fifth (19 percent) of the state's \$1 trillion GDP in 2020 (BEA 2020). We start by documenting some new stylized facts about wetlands trading. First, we find considerable trade, with more than \$1.1 billion of transactions in regional markets from 1995 to 2018. Second, we show that this industry is highly concentrated, with fewer than three wetland banks trading in an average market. Third, we find evidence of spatial reallocation of wetlands away from densely populated flood hazard areas into peripheral zones, consistent with private gains from trade as well as adverse selection in terms of local flood protection.

We then use observed offset trades, prices, and production to measure the private gains from trade and estimate a model to predict equilibrium wetlands reallocation and environmental outcomes under alternative market designs. The empirical strategy proceeds in three steps. First, we estimate demand for wetland offsets using transaction-level data on the location, price, and quantity of offset purchases over time. We build several price instruments from cost shifters of offset supply based on our understanding of the industry, such as variation in offsets issued to historical incumbents based on fixed production schedules and variation in public wetlands that affect feasibility of production.

Second, we estimate a model of industry dynamics of offset supply, using (i) administrative data on the set of operating wetland producers and (ii) maps that indicate the location of entrants. Our strategy for identifying the cost structure for this industry follows in the tradition of Bajari, Benkard, and Levin (2007) and Pakes, Ostrovsky, and Berry (2007) to leverage equilibrium conditions for firm behavior. We use observed offset production over time to directly estimate wetland production schedules as functions of fixed bank characteristics. To account for offset storage, we

¹ Wetlands comprise 6 percent of land worldwide and 12 percent of the terrestrial carbon stock (Erwin 2009), but their global extent declined by 35 percent between 1970 and 2015 (Ramsar Convention 2018).

characterize trading as an optimal inventory problem. We then combine estimates of offset demand with optimality conditions for entry, which allow us to obtain expected incumbent profits. We then estimate conditional entry cost distributions to rationalize observed entry decisions as solutions to each producer's dynamic optimization problem as in Bajari, Benkard, and Levin (2007). In particular, we obtain conditional entry cost distributions that depend on local characteristics that affect the feasibility of wetlands restoration across space.

To avoid the curse of dimensionality in estimating entry costs, we approximate strategic entry and trading decisions with functions that depend on a subset of rivals' characteristics, following work such as Ryan (2012); Ifrach and Weintraub (2017); and Gowrisankaran, Langer, and Zhang (2025). To circumvent the curse of dimensionality in counterfactuals, we follow the approach of Rafey (2023), who obtains realized gains from trade by integrating estimated value functions over observed trade flows, avoiding the need to calculate a new equilibrium. In addition, for our Pigouvian counterfactuals, we hold entry fixed, as well as observed bank-level trades unless they violate individual rationality constraints calculated from the counterfactual demand curves. We then obtain counterfactual developer values and entry costs by integrating the estimated demand curves and an aggregate cost function over the counterfactual bank trades. These restrictions enable us to report approximate market outcomes (private gains from trade, externalities, and total surplus) using the estimated model primitives, observed offset trades, and a set of computationally tractable constraints.

Third, to analyze environmental consequences of wetlands reallocation under the current market design, we estimate wetlands' local values for flood protection, a major hydrological outcome not currently incorporated into existing offset trading rules. In Florida, approximately \$700 billion of assets lie in a 100-year flood zone (Wing et al. 2018). Moreover, new empirical research suggests that the value of these local flood protection benefits may be considerable (Brody, Highfield, and Blessing 2015; Sun and Carson 2020). We estimate our local flood protection functions using detailed historical land use and flood insurance claims data. This allows us to evaluate the quality of newly produced offsets relative to direct conservation.

Our main empirical findings are threefold. First, we find substantial private gains from trade, reflecting the significant differences between the opportunity cost of developing marginal wetlands and the entry costs of wetland mitigation banks. Second, we find that by shifting wetlands away from places most vulnerable to flood risk, the market increased total flood damages, though these outcomes are highly heterogeneous across space. Third, we show that augmenting the current market design with Pigouvian taxes proportional to local flood risk can eliminate almost 80 percent of flood damages while preserving more than two-thirds of the private gains from trade. A uniform development tax also lowers total flood damage but leads to lower private surplus and significantly greater flood damages than the differentiated Pigouvian prescription.

Contributions to the Literature.—This paper makes three primary contributions. First, we provide an empirical framework for environmental market design in regulated conservation offsets. Methodologically, we build on both the literature that seeks to value the gains from trade under market-based reallocation relative to less

flexible environmental or energy regulations (e.g., Carlson et al. 2000; Borenstein, Bushnell, and Wolak 2002; Rafey 2023), as well as the literature on second-best pricing of heterogeneous externalities (Diamond 1973), which, in environmental economics, often emphasizes the dangers of environmental markets in second-best contexts where pollution occurs at finer gradations than policy instruments (e.g., Muller and Mendelsohn 2009; Fowlie, Reguant, and Ryan 2016; Fowlie and Muller 2019).

Second, we augment existing models of land use and conservation with landowners' restoration activities that produce offsets. Static models of long-run conservation and land use, such as Stavins and Jaffe (1990); Souza-Rodrigues (2019); and Assunção et al. (2023), as well as recent models of dynamic land use (e.g., Scott 2013; Hsiao 2026), rule out the use of land to supply new environmental protection. Here, we specify and estimate the production technology for new restoration projects, derive equilibrium outcomes for the concentrated markets that arise from the fixed costs and time-to-build of these technologies, and endogenize landowners' opportunity costs of meeting a given conservation objective through the offset market. Our empirical findings show how private costs of land-use restrictions (e.g., Saiz 2010; Turner, Haughwout, and van der Klaauw 2014), and wetland permitting specifically (Silverstein 1994; Keiser et al. 2022), can fall over time. Several far-reaching judicial decisions have relied on the general assumption that wetland permitting schemes are unduly burdensome for landowners, with the US Supreme Court repeatedly citing a seminal economic study from more than two decades ago by Sunding and Zilberman (2002) (cited in the first paragraph of *Rapanos v. United States* (2006, §1A), the second paragraph of *USACE v. Hawkes* (2016, §1A), and again in *Sackett v. EPA* (2023, §1A)). Our paper shows that this regulatory burden is neither inevitable nor invariant to the regulatory environment—the flexibility that offsets provide to landowners substantially lowered private compliance costs.

Third, we contribute to a growing literature on wetlands and hydrological outcomes. Our focus on the imperfect substitutability between original wetlands and new wetland banks in terms of flooding follows Rafey (2020), which built on work suggesting important interactions between conservation and floods (Kousky and Walls 2014) and connecting land-use data with flood outcomes (Brody, Highfield, and Blessing 2015; Sun and Carson 2020). Of particular note is Brody, Highfield, and Blessing (2015), who are the first to use land cover data to relate changes in wetland extent and flood insurance claims, as well as subsequent work by Taylor and Druckenmiller (2022) that relies on a similar dataset and empirical strategy as Brody, Highfield, and Blessing (2015). Like these papers, our work emphasizes the spillovers created by wetlands that protect existing property, and our research design relies on detailed hydrological and historical data.

We build on this prior work on wetland externalities in two ways. One, we connect our estimates of local wetland values directly to the economics of marginal wetland conservation and restoration. This allows us to estimate the effects of regional wetland markets, quantify their cost savings and flood externalities, and assess the welfare consequences of including flood externalities in the design of these markets. This paper is the first economic analysis to attain these objectives. Two, we improve the precision of wetland flood protection functions by (i) using a nonlinear model that more closely fits the data on flood damages and (ii) focusing

on spillovers to properties built prior to the market to reduce bias. We find wetlands deliver policy-relevant spillovers in some, but not all, places. Our findings differ considerably from recent US-wide estimates of such spillovers in Taylor and Druckenmiller (2022), which, when applied to Florida, exceed our flood protection estimates by more than an order of magnitude. This discrepancy indicates that actual policy evaluation requires carefully tailored approaches to estimating marginal wetland flood protection functions, using research designs that compare similar places with and without marginal wetlands.

Outline.—The rest of the paper is organized as follows. Section I provides background on the legal framework that governs activities that destroy, conserve, and restore wetlands, as well as motivating evidence for the sources of gains from trade and adverse environmental outcomes. Section II specifies a model of equilibrium supply and demand for wetland offsets, and Section III describes the empirical strategy and benchmark estimates. Section IV evaluates private gains from trade, local flood outcomes, and some counterfactual market designs to internalize flood risk; Section V concludes.

I. Background and Data

A. Basics of Wetlands and Offsets

Wetlands deliver an array of local public goods, but wetland conservation entails private costs. Wetlands consist of marshes or swamps and, in the continental United States, cover more land than the state of California (*Rapanos v. United States* 2006). Their multifarious environmental services are difficult to value and rarely priced.² At the same time, their conservation precludes alternative land uses and therefore can entail substantial economic cost, often borne by landowners whose property includes wetlands.

Activities that risk degrading local wetlands have been regulated in the United States since the 1972 Clean Water Act. Section 404 of the Clean Water Act prohibits economic activity that risks “significantly degrading” existing wetlands. This prohibition has been taken as a mandate to conserve an aggregate stock of ecological and hydrological functions delivered by wetlands. Under this “No Net Loss” principle, wetland degradation can occur if it is accompanied by approved actions that “offset” the degradation (EPA and Department of the Army 1990).

The first iteration of No Net Loss was prescriptive and did not involve trade. Land developers on existing wetlands were typically either denied permits or required to implement mitigation activities on-site (Salzman and Ruhl 2006), though in some cases, developers paid local “in lieu fees.” This nonmarket approach was heavily criticized by private landowners and environmental groups alike. Land developers argued that the requirements were unduly burdensome (Sunding and Zilberman 2002), while environmental stakeholders argued that on-site mitigation activities did not compensate fully for wetland loss (Ruhl, Salzman, and Goodman 2009).

²For example, wetlands can purify water, enable recreation, and sustain wildlife such as herons, alligators, and manatees, depending on attributes like location, age, maturity, and salinity.

Tradable offsets arose in response to these concerns. Rather than requiring land developers whose land included wetlands to undertake on-site mitigation actions or prohibiting development outright, landowners could buy offsets from wetland restoration projects, known as “wetland mitigation banks.” These projects commit land to the public trust and engage in varied conservation activities to restore degraded wetlands or create new ones (e.g., converting farmland back to its natural state Erwin (2009)).

Our empirical analysis focuses on offsets required by Florida state law for wetlands protected under the 1972 Florida Water Resources Act. With the greatest share of wetland cover of any state in the continental United States, and rapid population growth and real estate development over the last three decades, Florida is a litmus test for wetland mitigation banking. The Florida laws governing wetland banking date from February 1994 (FS §373.4135). Our focus on Florida rather than federal wetlands is motivated by two considerations, discussed in more detail in Supplemental Appendix B.5. First, Florida jurisdiction encompasses all wetlands in Florida, including those regulated under the federal Clean Water Act, as well as wetlands outside of federal jurisdiction, such as those not connected to the Atlantic Ocean or the Gulf of Mexico by navigable waters. Second, the jurisdictional boundaries of the Clean Water Act have shifted over time in response to legal and administrative changes (Keiser et al. 2022), whereas Florida jurisdictions have remained stable during our study period.

B. *Trading Rules*

Regulators enable and oversee several crucial aspects of the certification and trade of environmental offsets to enforce No Net Loss. Importantly, the regulator has permitting authority: Land developers must obtain approval before either developing or restoring wetlands. To this end, the regulator defines exchange values between restored and existing wetlands through on-site assessments and a uniform assessment method.³ Although assessments incorporate diverse criteria related to biodiversity and ecological integrity, they do not directly account for the flood protection that wetlands can provide to the surrounding built environment.

For development on protected wetlands, the regulator evaluates the development’s adverse effect on regional wetland functionality, then specifies the offsets the developer needs to purchase in order to proceed. A developer who buys offsets from a bank is limited to purchasing offsets from a bank operating within the same hydrological region (Figure 1, panel A). These market boundaries, known as wetland mitigation bank service areas, approximate hydrological regions and extend far beyond the local project site.

For wetland mitigation banks, the regulator requires an environmental audit, a set of proposed restoration activities, and a detailed implementation schedule and cost budget. In addition, committing land to a wetland bank requires a permanent conservation

³To define equivalent units across diverse wetlands, regulators use the “uniform mitigation assessment method,” which defines exchange ratios across wetland attributes to deliver a scalar measure of wetland value. The method captures the “ecological and hydrological functions” a wetland delivers to the surrounding region (Florida State Legislature 2019, §373.4136(1)); for example, a bank might deliver ecological functions by planting trees or removing invasive species, and hydrological functions by building dams or canals.

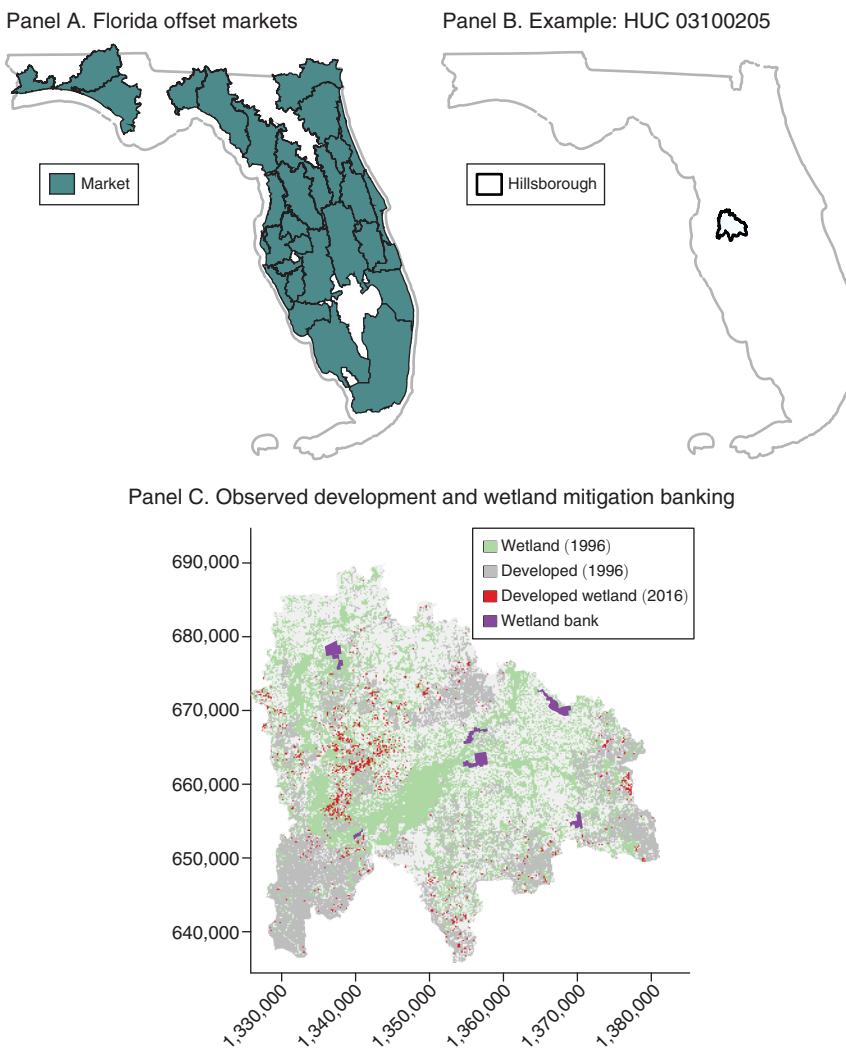


FIGURE 1. LOCATIONS OF WETLAND DEVELOPMENT AND RESTORATION

Notes: An example of our data on land use and wetland offsets within an offsets market. Initial wetland (green) and developed (gray) pixels in 1996, new development on wetlands from 1996 to 2016 (red), and wetland bank parcels established by 2018 (purple). Supplemental Appendix Figures A9.1–30 replicate this map for every market in our study. Table 7 reports average differences between all watersheds, watersheds with development (red pixels), and watersheds with wetland mitigation bank sites (purple pixels).

easement, ruling out alternative future land uses. Each project’s total lifetime output reflects the regulator’s assessment of its contribution to wetland functionality. Total lifetime production is specified at the time of entry, with offsets released gradually as wetlands regenerate and the regulator verifies that the bank attains its restoration goals. Banks can, and do, hold offsets in reserve to sell in future periods.

For offset trading, the regulator maintains a ledger that tracks the creation and retirement of wetland offsets. The regulator issues offset credits to wetland banks, verifies that buyers obtain sufficient offsets to compensate for their development,

and deletes the corresponding offsets from the bank's balance. While the ledger is centralized and maintained by the regulator, offset trades between wetland banks and land developers occur bilaterally. Such over-the-counter trades are typically brokered through private intermediaries. This decentralization makes the exact market mechanism unknown. Actual trading may exhibit a variety of imperfectly competitive features.⁴

C. Data Sources

We develop a new dataset to track wetlands, development, and offsets across Florida from 1995 to 2018. Our work draws on several new primary sources summarized in Supplemental Appendix Table A1 and detailed in Supplemental Appendix A. Here, we briefly describe the novel aspects of our data, emphasizing how these sources reveal (i) the timing, origin, destination, and volume of offset trade flows; (ii) prices for offset trades; (iii) land ownership, assessed values, and prices; (iv) flood risks and damages; and (v) wetland location and extent.

First, we track offset trading with administrative data on environmental permits and offsets from the Florida Department of Environmental Protection (FDEP) and regional water management districts (FDEP 2014b,c; SFWMD 2016; FDEP 2022b; SFWMD 2022; SJRWMD 2022b; SWFWMD 2022a). These agencies regulate the creation and sale of offsets and licenses for wetland restoration and conversion. From their records, we assemble a comprehensive ledger of the location, timing, and quantity of all state wetland offset transactions in Florida from 1995 to 2018. In addition, we obtain detailed producer-level data for every wetland mitigation bank operating over this period (FDEP 2022a; SFWMD 2022; SJRWMD 2022a; SWFWMD 2022b). Entry requires certification from either FDEP or water management districts, who maintain contracts with every wetland bank in Florida. These contracts include maps of the bank site, the date of the initial contract, and details on the offset release schedule over time. Many contracts also include reported restoration costs, which we use to corroborate our estimates.

Second, we obtain prices for wetland offset transactions from market participants. Our main source is a nondisclosure agreement with a major private broker (Mitigation Marketing, LLC 2018). We supplement the data on these private transactions with Freedom of Information Act requests to county officials (Brevard County 2018; Collier County 2018; Lee County 2019) and the Florida Department of Transportation (FDOT 2018) for government offset purchases. While transaction prices are not reported to the regulator, our final data include the majority of trades and nearly the entire period (1998–2018).

Third, we construct maps that track evolving environmental characteristics of coastal land to measure wetland location, extent, and quality. These land cover maps are derived from satellite and aerial data in the National Oceanic and Atmospheric Administration's (NOAA) Coastal Change Analysis Project (C-CAP) and cover all of Florida at a 30m × 30m resolution in 1996, 2001, 2006, 2011, and 2016

⁴For example, offset procurement by the Florida Department of Transportation and many local governments involve sealed-bid auctions. Private sales, by contrast, involve bilateral negotiations, and, at the same time, intermediaries typically post price lists for their prospective clients.

(NOAA 2020). These data contain more than 194 million pixels for each of 5 periods, 136 million of which are contained in offset trading zones, giving us an unprecedented view of the evolution of land use in Florida.⁵

Fourth, we use maps of land ownership to delineate between private and public wetlands. We use boundaries of all land owned by local, state, and federal entities at baseline (1995) and Florida conservation purchases from 1990 to 2020 under the Preservation 2000 and Florida Forever programs (FDEP 2014a, 2018). We also use annual zip code–level home values from Zillow (2020) for 1998–2020 and US Census Bureau (2009–2021, 2000) data on population, income, housing units, and home values for 2000 and 2007–2019.

Fifth, we collate local flood data from the Federal Emergency Management Agency (FEMA). Our primary measure of economic damages uses administrative data from FEMA, which administers virtually all flood insurance contracts and claims. We use recently redacted, publicly available data on the universe of flood insurance claims and policies from 1978 to 2020, which include the claim location, date, and amount (FEMA 2023), as well as data obtained through a FOIA request that includes total policies held from 1975 to 2018. In addition, we calculate local measures of inherent flood risk using flood zone designations from the National Flood Hazard Layer (NFHL) from FEMA (2020). The NFHL is based on topographical and hydrological modeling. These detailed maps of flood risk are used to price flood insurance at the city-block level and capture all locations, whether or not they have purchased insurance.

We then match the diverse spatial and temporal scales of the microdata to build a hydrologically consistent panel as described in Supplemental Appendix B. Specifically, we use hydrological boundaries from the United States Geological Service (USGS 2022) to produce a consistent panel of local watersheds and markets across time that aligns with both hydrological realities and market boundaries. Local watersheds are typically about 24,000 acres (40 square miles). Florida contains 1,378 such watersheds, 1,004 of which lie within offset markets (Supplemental Appendix Figure A1).

D. Descriptive Evidence

We now use our data to outline some facts about (i) initial wetland extent and land ownership, (ii) spatial patterns in development and wetland restoration, (iii) offset releases and sales, (iv) market structure, and (v) trade outcomes.

First, many privately owned wetlands exist throughout Florida at baseline. Table 1 shows that 36.4 percent of the 136,302,645 pixels in our dataset are initially a wetland, with the average watershed containing 10,818 acres of wetland (33.2 percent of its area). Many, but not all of these wetlands will be prospective sites for development, depending on the initial ownership of the land; in our data, private wetlands account for 99.2 percent of wetland pixels developed over the sample period. Local, state, or federal entities own 12.5 percent of an average watershed's area in 1995 and

⁵Relative to more general land cover datasets, C-CAP is tailored to study coastal systems in the Gulf Coast (6 of its 25 land use categories are wetland subtypes; see Supplemental Appendix D) over time, tracking actual wetland transitions with high levels of accuracy (McCombs et al. 2016).

TABLE 1—NEW DATA ON WETLAND OFFSETS IN FLORIDA

	Observations	Avg	SD	q25	q50	q75
Initial wetlands ^a (pct/pixels)	136,302,645	36.5	48.1	0	0	100
Initial wetlands (pct/watershed)	1,004	34.0	20.1	20.0	30.8	44.0
Initial public land (pct/watershed)	1,004	12.5	21.8	0	2.1	14.7
Wetlands development and restoration						
Wetlands developed, 1996–2016 (acres)	1,004	207.5	483.4	2.4	16.3	186.7
Private wetlands developed	1,002	206.2	481.3	2.0	15.9	186.1
$\mathbb{P}(\text{develop} \text{wet}) \times 100^b$	1,000	3.7	7.3	0.04	0.3	3.9
With wetland bank ^c	96	0.8	1.1	0.1	0.2	1.0
With high development ^d	179	15.1	10.1	7.0	14.2	20.8
Initial wetlands ('000 acres)	1,004	11.0	30.6	4.1	7.2	11.6
With wetland bank	96	23.1	76.1	7.6	10.5	15.9
With high development	179	10.1	13.8	3.5	7.0	11.2
Initial public wetlands ('000 acres)	1,004	3.3	25.0	0	0.3	1.9
With wetland bank	96	13.7	74.1	0	0.6	3.9
With high development	179	2.0	6.3	0.1	0.4	1.6
Initial private wetlands ('000 acres)	1,004	7.7	13.8	3.2	5.7	9.3
With wetland bank	96	9.4	5.5	5.8	8.4	11.2
With high development	179	8.1	8.8	3.3	6.0	9.6
Initial developed land (pct)	1,004	13.8	19.1	1.8	4.7	18.2
With wetland bank	96	5.7	6.6	1.2	3.0	8.0
With high development	179	37.6	23.0	19.1	32.5	55.7
Offset credit production and sales						
Bank entry year	107	2008.1	7.5	2003	2009	2014.5
Bank size (acres/bank)	107	1,866.1	2,680.0	428.5	1,049	2,157.5
Bank size (credits/bank)	106	410.0	566.1	85.2	203	521.8
1 (credits released) per bank per year	1,209	0.3	0.4	0	0	1
Credits released (pct/bank/year)	343	15.3	16.2	5	10.0	20.0
Acres per credit	106	5.9	4.5	3.1	5.1	6.9
Acre wetland developed per credit sold	5,512	8.8	2.8	8.1	8.2	11.6
Annual sales (credits/bank-year)	981	15.5	31.4	0	1.8	15.4
Bank reserves (pct/bank-year)	967	51.8	33.6	18.3	54.7	82.0
Market structure						
Area ('000 acres/market)	30	1,153.3	800.4	516.8	863.4	1,520.4
Area (watersheds/market)	30	33.5	18.0	18.2	27.5	49
Entry (market-year)	780	11.7	32.1	0	0	0
Number of banks (market-year)	530	2.6	2.3	1	2	3
Offset price ('000\$/credit), all transactions	1,432	87.5	61.7	38.6	63.4	137.2
Offset price ('000\$/credit/market/year)	151	98.8	50.5	62.0	93.9	127.2
Flood risks						
Flood zone (pct/watershed)	1,004	41.7	23.8	23.9	37.3	56.1
Zone V (storm surge) (pct)	1,004	2.4	9.8	0	0	0
Zone A (100-yr) (pct)	1,004	39.4	22.5	23.0	35.6	51.7
Flood insurance claims ^e ('000\$/claim)	188,368	31.3	71.1	3.3	10.5	32.9
Flood claims, pre-1996 ('000\$/yr/watershed)	1,004	219.9	1,387.4	0	0.2	10.1
With wetland bank	96	314.6	2,345.2	0	0.1	4.7
With high development	179	412.7	1,552.4	0.8	10.7	99.4
Flood claims, post-2015 ('000\$/yr/watershed)	1,004	335.2	1,414.6	0.003	4.9	71.5
With wetland bank	96	161.5	541.2	0.7	7.1	49.6
With high development	179	798.2	2,445.0	24.1	97.3	361.1

Notes: Descriptive statistics for Florida, 1995–2020. Supplemental Appendix Tables A2 and A3 contain additional data.

^aInitial measures correspond to 1996 values.

^b $\mathbb{P}(\text{develop}|\text{wet})$ defined as the within-pixel probability that a wetland pixel in 1996 becomes developed in 2016.

^cWatersheds with at least 100 acres of a wetland bank site and fewer than 250 acres of developed wetlands.

^dHigh-development watersheds defined as those with greater than 250 acres of developed wetland from 1996 to 2016 and fewer than 100 acres of a wetland bank site.

^eAll flood insurance claims from 1985 to 2020.

2.1 percent of the median watershed. Wetlands are more likely to be publicly owned than other types of land, but more than two-thirds of all wetlands in Florida, or about 7,400 acres per watershed, are privately owned.

Second, our spatial data reveal systematic patterns in development and wetland restoration.⁶ Figure 1, panel C illustrates the typical pattern of reallocation using within-pixel data in a representative market. Wetland development (red pixels) occurs nearer historical development (dark gray), while wetland bank project sites (dark blue) are fewer, closer to historical wetlands (green), and farther from developed areas. Similar core-periphery patterns are apparent in the other 29 markets that we study, depicted in Supplemental Appendix Figures A9.1–30. To quantify these patterns, Table 1 compares watersheds that contain wetland banks to watersheds with substantial wetland development.⁷ Most development occurs in places with greater initial development density: 32.5 percent of the area of the median high-development watershed starts as developed, vastly exceeding the median watershed's 4.7 percent or the median wetland bank watershed's 3.0 percent. Wetland development also occurs frequently alongside other land development, with a correlation between development on wetland pixels and contemporaneous development of other pixels in a watershed of 0.656.

Wetland banks, in contrast, enter in watersheds with more initial public wetlands (13,700 acres) than the average Florida watershed (3,300 acres). This pattern is consistent with regulatory incentives that award additional offsets to banks to restore existing wetlands near existing conservation land,⁸ hydrological advantages of restoring wetlands nearby other wetlands, and lower land values in peripheral zones with significant conservation area.

Large differences between wetland development and restoration sites also exist in terms of flooding. High-development watersheds have greater insured value at baseline than wetland bank watersheds (\$18.8 million compared with \$10.1 million); both groups have more than the average watershed in the sample (\$7.2 million) because offset trading occurs in places with disproportionate flood risk. High-development and wetland bank watersheds have similar average historical flood insurance claims (\$413,000/year versus \$314,000/year), but this difference is not statistically significant due to the immense dispersion of these distributions, which have coefficients of variation greater than 5. In the post-period, average flood insurance claims double in watersheds with high development, to \$800,000 per year by 2016–2020. In contrast, average claims in wetland bank watersheds decline in real terms from \$315,000 to \$160,000 (2020 US dollars) per year, and the

⁶To determine where offsets are produced, we match wetland bank locations to watersheds. To determine where offsets are bought, we use the conversion of private wetlands into developed land. In our structural analysis, we analyze outcomes only for wetlands converted in places where we observe offsets trade; this corresponds to about two-thirds of all wetland development because some markets begin after 2000.

⁷We define high-development watersheds as those with at least 250 acres of wetlands developed, about 10 percent of the watersheds. While the average watershed converts 207.5 wetland acres over the study period, the median watershed sees little development (16.3 acres), whereas the seventy-fifth percentile watershed sees 186.7 acres converted; initial wetlands are developed with probability 0.037 in the average watershed but probability 0.57 in the watershed with the greatest share of wetlands developed.

⁸"Mitigation banks and off-site regional mitigation should emphasize the restoration and enhancement of degraded ecosystems and the preservation of uplands and wetlands as intact ecosystems rather than alteration of landscapes to create wetlands" (Florida State Legislature 2019, §373.4135).

median bank watershed sees fewer than one-tenth of the flood claims of the median high-development watershed.

Third, our production data show that banks produce large quantities of offsets relative to the size of their markets and of offset trades. The median bank produces about 200 offsets over its lifetime (or 410 on average), with an interquartile range of [85, 520] offsets, while the median developer purchases only 1.1 offsets (or 4.1 offsets on average).⁹ The scale economies for wetland banks reflect the large parcel areas required to redirect water flows, as well as rules for banking that reward wetland contiguity. Production increases with the total area of the wetland bank project site; on average, the ratio of acres to offsets is about 5.9 acres, ranging between 3.1–6.9 across water management districts. Banks also take time to build, reflecting the need to verify environmental improvements over time. Table 1 shows that the regulator typically releases 15 percent of a wetland mitigation bank's offsets once every 3 years, or an average of $1/0.055 \approx 18.2$ years to build the entire project.

Fourth, we find that Florida contains many distinct offset markets because banks can only trade offsets within their regional service areas and enter relatively infrequently. Figure 1 depicts the market boundaries we use, defined using hydrological regions after some adjustments to correct for partial and overlapping service areas, as discussed in Supplemental Appendix B. On average, a market covers 1.15 million acres, or 33.5 watersheds, ranging from 11 to 70 watersheds. Wetland banks enter in 11.7 percent of market-years, such that the median market in the median year (2006–2007) contained 1 bank or 2.2 on average, rising to 2 incumbents and an average of 3.7 firms by 2018. Supply is often concentrated among a few banks; the average bank owns 26.1 percent of its market's total production potential and 37.4 percent of its market's total unsold offsets. The latter reflects the fact that banks rarely sell their offsets immediately but rather hold positive reserves; the median bank holds 52 percent of its offsets in reserve, with an interquartile range of 18–82 percent. We interpret this concentration in offset supply as reflecting economies of scale and production delays discussed above, as well as strategic factors that interact with the regional restrictions on trade.

Fifth, our ledger and bank contract data directly reveal some realized trade outcomes. Valued at average annual prices, cumulative offset sales totaled \$1.1 billion from 1995 to 2018 (in 2020 US dollars), making offsets an increasingly central feature of Florida wetland management since 1995. Offset sales increase over time, growing annually by an average of 9.8 percent, reflecting demographic shifts driving new development in Florida as well as the transition to market-based wetland conservation after the introduction of wetland bank rules in 1994. In addition, offset prices considerably exceed observed components of bank costs. Average real prices over the full sample are about \$88,000 per offset. Total costs of banks for which we see cost data average \$5.3 million, or about \$24,000 per offset. Restoration costs put in escrow, which we observe for nearly two-thirds of banks, average \$7,000/offset, while land values obtained from the last reported transaction price average \$19,000/offset (\$9,000/offset). Variation across banks appears to reflect local land

⁹Measuring the location of these projects is more difficult than for wetland banks; we observe the quantity and timing of offset sales by each bank, and therefore in each market, but we do not observe every parcel that purchases offsets. See Supplemental Appendix C.1 for a discussion.

prices as well as natural features that determine the costs and feasibility of restoration across markets.¹⁰

Taken together, our data indicate substantial trade flows between wetland banks and wetland developers that locate in quite different places even within relatively small regional markets. These patterns indicate both large prospective private gains from trade—because marginal wetlands used for restoration or converted into development differ meaningfully along observables—as well as the possibility of first-order changes in hydrological externalities like flooding. However, transaction volumes and prices are not sufficient statistics for the gains from trade.¹¹ Further, selection into mitigation banking precludes the direct use of cost data from our contracts for counterfactuals, which require the unconditional cost distribution of all prospective banks, not just those that entered. Entry also involves costs not observed from contracts—such as permitting costs—and entry incentives further depend on the value of equilibrium trading over time. Moreover, evaluating offsets' effects on other outcomes like flood externalities requires identifying the relative effects of marginal wetlands and new wetland banks on these outcomes. The empirical model of decentralized trade in environmental offsets below is designed to tackle these issues.

II. A Model of Conservation, Destruction, and Restoration

We now specify an empirical model of regulated environmental offsets. Wetlands distributed across space can be conserved, developed, or restored over time (Section IIA). A regulator issues permits to ensure that offsets satisfy its various conservation objectives (Section IIB). Small land developers take offset prices as given and obtain payoffs from developing existing wetlands (Section IIC). Large producers restore wetlands to obtain offsets from the regulator, which they can sell to land developers over time. These producers incur fixed entry costs, zero marginal costs, and take time to build (Section IID). Incumbents simultaneously choose sales in each period in a Markov perfect equilibrium (Section IIE). In this setting, entry follows a cutoff rule, and dynamic trading strategies can be characterized as an optimal inventory problem (Section IIF).

A. The Conservation Problem

A large hydrological region, or “market,” consists of a map of a continuum of locations indexed by $i \in [0, 1]$, which we partition into a finite set of local watersheds, indexed by h . Within a market, the distribution of wetlands at time t is given by $\{w_{it}\}_i$, with $w_{it} = 1$ when i contains a wetland at t , and $w_{it} = 0$ otherwise. Time is discrete, the horizon is infinite, and all agents discount future periods with a factor $\delta < 1$.

¹⁰For example, restoration costs are lower in Northern Florida (e.g., about \$9,000/offset in Altamaha–St. Mary’s) than Gulf Coast markets (e.g., \$16,000 in Peace–Tampa). Similarly, land costs are higher in Southern Florida (\$12,000/offset) and the Gulf (e.g., \$12,600/offset in Peace–Tampa) than northern markets (e.g., \$5,700/offset in Altamaha–St. Mary’s).

¹¹Inframarginal buyers may have values significantly greater than market prices, while imperfect competition may allow banks to charge prices above their costs.

Wetland conservation, development, and restoration occur over time. Each of these processes correspond to a different state transition between t and $t + 1$. First, existing wetlands can be conserved; that is, $w_{it} = w_{i,t+1} = 1$. Second, locations with wetlands can be developed into nonwetland property and sold; that is, $w_{it} = 1$ and $w_{i,t+1} = 0$. Third, land without wetlands can be restored into wetlands; that is, $w_{it} = 0$ and $w_{i,t+1} = 1$.

Wetlands have social and private values. The social value of wetlands arises through their diverse attributes, $v_i \in \mathcal{V}$ for each i , where \mathcal{V} is some set of attributes. The private costs and benefits of wetland conservation, development, and restoration accrue to landowners and wetland restoration firms. Given the private payoffs of wetland conservation, development, and restoration, in each period t , landowners will decide land use for $t + 1$ and incur costs of land-use change. Importantly, not all land-use decisions are reversible. We model restoration (a transition from $w_{it} = 0$ to $w_{i,t+1} = 1$) as an absorbing state, given that wetland banking requires a permanent transfer of land ownership into the public trust (a conservation easement). Similarly, we model development (a transition from $w_{it} = 1$ to $w_{i,t+1} = 0$) as an absorbing state. This is because, in the three decades spanned by our data, wetlands converted to development almost never transition back to wetlands.

B. Offset Market Design

The regulator enforces No Net Loss in “wetland value,” a function that maps wetland attributes, $v_i \in \mathcal{V}$, into a number of offsets, $\tilde{v}_i \in \mathbb{R}$. The regulator’s measure of aggregate wetland value is

$$(1) \quad \tilde{v}(w_t) = \int_0^1 w_{it} \tilde{v}_i di,$$

and No Net Loss requires that the distribution of wetlands $\{w_{it}\}_i$ delivers at least as much value in each period t as in the initial period; that is,

$$(2) \quad \tilde{v}(w_t) \geq \tilde{v}(w_0) \text{ for all } t > 0.$$

In practice, the regulator enforces (2) by certifying sufficient cumulative wetland restoration to offset cumulative wetland destruction.

Offset trades to satisfy (2) involve two types of participants. First, owners of wetlands with development potential seek approval from the regulator to build. The regulator inspects each such location i to determine its environmental value, \tilde{v}_i , and then approves the project when the developer proves that they have purchased \tilde{v}_i offsets. Second, prospective mitigation bank entrants, indexed by f , propose restoration to the regulator. The regulator inspects each location f to determine \tilde{v}_f and the bank decides whether to enter and incur entry costs. The regulator monitors and verifies restoration activities and issues \tilde{v}_f offsets over time as the restoration succeeds. These offsets can be held by the incumbent wetland bank and traded in any future period.

Importantly, the irreversibility of both development and restoration simplifies the dynamic land-use problem in our setting by allowing us to separate private land into two types based on the initial conditions: first, prospective “developers” with

$w_{i0} = w_{it} = 1$, who decide in each period whether or not to develop their wetland into something with greater private value; second, prospective “wetland mitigation banks,” with $w_{i0} = w_{it} = 0$, who decide whether or not to enter. We analyze each type’s decision in the next two sections.

C. Demand for Offsets

Developers must purchase offsets to build on their wetlands. We assume a competitive market for private land development with a continuum of landowners, indexed by i , populating a finite collection of watersheds h . Landowners $i \in h$ who develop on a wetland at t (i.e., i such that $w_{i,t+1} < w_{it} = 1$) obtain a private value of development given by

$$u(X_{ht}, \xi_{ht}; \theta) + \epsilon_{it1} = \theta' X_{ht} + \xi_{ht} + \epsilon_{it1},$$

which has two parts. First, an ex ante value of wetland development, $u(X_{ht}, \xi_{ht}; \theta)$, which depends on observed local characteristics X_{ht} (such as development density, demographics, hydrological region, and local flood risk), unobserved local characteristics ξ_{ht} , and a vector of preference parameters θ . For example, this ex ante value can correspond to the discounted stream of rental income from developed land or expected profits from agricultural production for land used to grow crops, net of the construction or future planting costs. Second, choice-specific, idiosyncratic costs of development and nondevelopment, ϵ_{it1} and ϵ_{it0} , independently and identically drawn over i and t according to a Type 1 Extreme Value (T1EV) distribution.

Without regulation, the ex ante private value for a landowner who develops on wetlands in h in period t is just $u(X_{ht}, \xi_{ht}; \theta)$, which determines the share of that watershed’s existing wetlands developed in a given period. However, in the market design of Section IB, developing on wetlands also requires offsets. If developer $i \in h$ can purchase offsets at a price P_t , then, given the regulator’s assessment \tilde{v}_h of i ’s watershed’s contribution to conservation priorities and a price sensitivity coefficient θ_p , i ’s relative value of destroying the wetland becomes

$$(3) \quad u(X_{ht}, \xi_{ht}; \theta) - \tilde{v}_h \theta_p P_t + \epsilon_{it1} - \epsilon_{it0}.$$

We assume that i destroys its wetland at t if and only if (3) exceeds zero. Aggregate demand for offsets at t at a price P_t and a regulatory rule \tilde{v}_h (offsets per acre in h) is then

$$(4) \quad Q_t(P_t, W_t, X_t, \xi_t, \tilde{v}; \theta) = \int_0^1 w_{i0} w_{it} \tilde{v}_h \mathbf{1}\{u(X_{ht}, \xi_{ht}; \theta) + \epsilon_{it1} - \epsilon_{it0} \geq \tilde{v}_h \theta_p P_t\} di$$

$$= \sum_h \tilde{v}_h W_{ht} \frac{e^{\theta' X_{ht} - \tilde{v}_h \theta_p P_t + \xi_{ht}}}{1 + e^{\theta' X_{ht} - \tilde{v}_h \theta_p P_t + \xi_{ht}}},$$

where the second line follows from the logit assumptions across local landowners. Aggregate demand in (4) reflects current shocks to local development payoffs, $(X_t, \xi_t) = \{X_{ht}, \xi_{ht}\}_h$, as well as the extent of private wetlands available for development, given by $W_t = \{W_{ht}\}_h$ across local watersheds h , with $W_{ht} \equiv \int_{i \in h} w_{i0} w_{it} di$ for each h .

The structure of the private landowner's decision above imposes some limitations on our analysis. First, while private wetland owners have the same average development payoffs within local watersheds—allowing for correlation across these development decisions—developers act independently from one another and take offset prices as given. These assumptions reflect the small size of these developers relative to one another and to the banks described in Section ID but rule out coordinated development schemes across many parcels. Second, while prospective developers capture the full value of new wetland development, they are otherwise myopic. That is, the decision rule in (3) rules out more complicated forward-looking strategies by developers that incorporate the option value of future development (e.g., as in Scott 2013). This restriction buys us considerable tractability on the demand side of our model but limits our analysis to the extent that individual developers delay development to obtain more favorable offset prices or choice-specific shocks.

Despite these restrictions, our model of wetland development captures some essential aspects of the economic setting. Aggregate market demand arises from many local watersheds, each with its own average utility, so our estimates of demand and consumer surplus capture variation across local watersheds in their revealed preference for developing wetlands, not only idiosyncratic logit shocks across landowners. Additionally, demand exhibits dynamics within watersheds and at the market level. Local stocks of potentially developable wetlands, W_{ht} , evolve endogenously with landowners' decisions. For example, greater development on wetlands today in a local watershed h will leave fewer prospective locations tomorrow, lowering W_{ht} and altering future demand for offsets. Furthermore, development on wetlands increases local development density, which itself affects the value of future development.

Over time, local demand also evolves with exogenous demand shifters. We assume these follow first-order Markov processes, that is, that the cumulative distribution function of (X_{t+1}, ξ_{t+1}) is some function $H_{X, \xi}(\cdot | X_t, \xi_t)$. This is without loss of generality; any finite-order Markov process admits a first-order representation under the appropriate extension of the state space. In Section III, however, we further restrict ξ_t to rule out persistence in unobserved and idiosyncratic watershed payoffs over time.

D. Supply of Offsets

We now turn to the choice problem for wetlands restoration, which—in contrast to dispersed development on wetlands—involves a few large restoration sites in each market. We model offset supply as an imperfectly competitive, dynamic oligopoly game with a finite set of noninfinitesimal potential producers, indexed by their location, $f \in \{1, 2, \dots, F\}$. Each production site f corresponds to a subset $I_f \subset [0, 1]$ of positive measure where restoration is feasible and $w_{i0} = 0$ for all $i \in I_f$. Production sites differ in natural suitability for restoration as well as intrinsic

production potential, \tilde{v}_f , which reflects various wetland services valued by the regulator, such as contiguity of the site with existing conservation land.

Entry.—In each period t , one potential entrant arrives at an unoccupied production site f at random, observes its potential environmental value \tilde{v}_f (denominated by the regulator in offsets), and then draws a private entry cost

$$(5) \quad \kappa_{ft} \sim G_t(\cdot \mid \tilde{v}_f, \mathcal{F}_t^c),$$

where G_t is a cumulative probability distribution conditional on \tilde{v}_f and observable local characteristics of the remaining production sites in the market, denoted by \mathcal{F}_t^c . The fixed cost captured by κ_{ft} includes permitting, restoration, and maintenance costs as well as the opportunity cost of nonwetland use. It may also include other aspects of operating the bank, such as intrinsic enjoyment of conservation. If the prospective entrant chooses to enter, the decision is irreversible, as discussed above. Otherwise, as in Doraszelski and Satterthwaite (2010), the prospective entrant disappears.

Production.—A bank produces offsets over time up to its total value, \tilde{v}_f . Because verification occurs gradually, the offset release schedule also depends on the bank’s age, T_{ft} . Specifically, in each period t , the regulator issues

$$(6) \quad b_{ft} = \mathcal{B}(T_{ft}, \tilde{v}_f),$$

offsets to each production site f . Offsets are issued until restoration is complete, that is, until $\sum_t b_{ft} = \tilde{v}_f$. Equation (6) allows for various time paths of offset release and also allows offsets’ release to occur stochastically but assumes that the restoration undertaken by the bank can be reasonably approximated with a known function of its land’s underlying characteristics, with capacity fixed in the initial contract and not revisable thereafter.

Trading.—Wetland banks obtain revenue by selling offsets to developers. At the start of each t , each incumbent f has a stock of available offsets $B_{ft} \geq 0$, certified but not yet sold. Each incumbent f can sell up to this constraint, $q_{ft} \leq B_{ft}$. Restoration costs are paid up front, so the marginal costs of producing and transacting offsets are zero.

Within each period, each firm f simultaneously chooses a quantity of offsets to trade, q_{ft} , which determines the price vector P_t via (4), and firm per period profits,

$$(7) \quad \Pi_{ft} = P_t' q_{ft}.$$

New wetland offsets, b_{ft} , are certified at the end of period t , and bank f ’s stock evolves to

$$(8) \quad B_{f,t+1} = b_{ft} + B_{ft} - q_{ft},$$

with the initial condition $B_{ft} = 0$ for all t prior to entry.

E. Information and Timing

We denote the market state vector at time t by

$$(9) \quad s_t = \left(W_t, X_t, \xi_t, \mathcal{F}_t^c, \left\{ \tilde{v}_f, B_{ft}, T_{ft} \right\}_{f \in \mathcal{F}_t} \right),$$

which consists of undeveloped private wetlands, $W_t = \{W_{ht}\}_h$, local characteristics (X_{ht}, ξ_{ht}) for each h , the remaining production sites \mathcal{F}_t^c , and the ages T_{ft} , offset balances B_{ft} , and capacities \tilde{v}_f for all incumbents $f \in \mathcal{F}_t$.

In each period t , all potential and current offset producers observe the market state, s_t . One prospective entrant $f \in \mathcal{F}_t^c$ then privately draws their fixed entry cost, $\kappa_{ft} \sim G_t(\cdot | \tilde{v}_f, \mathcal{F}_t^c)$, and decides whether to enter. Incumbents simultaneously choose their trading volumes, $\{q_{ft}\}_{f \in \mathcal{F}_t}$, which determines equilibrium offset prices via (4), and banks obtain profits. Finally, entry occurs, wetlands are developed, and the state updates to s_{t+1} .

F. Equilibrium

We focus on Markov perfect equilibria (MPE) (Ericson and Pakes 1995; Maskin and Tirole 2001) as formalized in Doraszelski and Satterthwaite (2010), restricting the strategies for each production site f to be anonymous, symmetric, and Markovian, so that they are given by functions

$$\sigma_f : (s_t, \kappa_{ft}) \mapsto (\text{enter}_{ft}, q_{ft}).$$

In an MPE, equilibrium profits within a period depend only on the wetlands available for private development, demand shocks, and incumbents' trading strategies and can be written as

$$\Pi_{ft} = \Pi(q_{ft}, s_t).$$

Firms maximize their expected discounted profits. The expected value of a wetland bank with offsets B and age T is

$$(10) \quad V(B, T, s_t, \tilde{v}_f) = \max_{q \in [0, B]} \Pi(q, s_t) + \delta \mathbb{E}_t[V(B - q + b_{ft}, T + 1, s_{t+1}, \tilde{v}_f)].$$

A bank's current trading decision affects its continuation value in two ways: first, directly, by depleting its future stock $B_{f,t+1}$; second, indirectly, through the state of undeveloped wetlands W_{t+1} , which affects future offset demand and entry incentives. We assume that the optimal trading decision at t , which maximizes (10), can be characterized by a function

$$(11) \quad q_{ft} = \mathcal{Q}(s_t, B_{ft}, T_{ft}, \tilde{v}_f)$$

of B_{ft} , T_{ft} , \tilde{v}_f , and s_t . Equation (11) assumes that \mathcal{Q} is a well-defined function—that is, that there is a unique equilibrium trading strategy at each state—but does not further specify conduct in the trading stage game.

All potential entrants use a common entry strategy that takes the form of a conditional cutoff rule: The pure strategy prescribes entry if and only if

$$(12) \quad \kappa_{ft} < V(0, 1, s_t, \tilde{v}_f).$$

This implies that the probability that f enters at t prior to its private draw of $\kappa_{ft} \sim G_t(\cdot | \tilde{v}_f, \mathcal{F}_t^c)$ is given by

$$(13) \quad \mathbb{P}(\text{enter}_{ft} | s_t) = G_t(V(0, 1, s_t, \tilde{v}_f) | \tilde{v}_f, \mathcal{F}_t^c),$$

which can be written as some function $\phi_t(s_t) \equiv G_t(V(0, 1, s_t, \tilde{v}_f) | \tilde{v}_f, \mathcal{F}_t^c)$.

The equilibrium in the environmental offsets market consists of entry and trading strategies $(\text{enter}_{ft}, q_{ft})_{t \geq 0}$ for all $f \in \{1, 2, \dots, F\}$, undeveloped private wetlands $(W_t)_{t \geq 0}$, and a path of offset prices $(P_t)_{t \geq 0}$, such that (i) entry satisfies (12) at all $t \geq 0$ for all $f \notin \mathcal{F}_t$; (ii) incumbents' trading strategies $(q_{ft})_{t \geq t'}$ solve (11) for all $f \in \mathcal{F}_{t'}$ and all t' ; (iii) private wetlands destruction Q_t solves (4) for every t ; and (iv) no net loss holds; that is, $\sum_{f \in \mathcal{F}_t} q_{ft} = Q_t$ for all t , as well as $\lim_{t \rightarrow \infty} \delta^t P_t' B_{ft} = 0$ for all f .

An MPE in symmetric pure strategies exists for this game by Doraszelski and Satterthwaite (2010, Proposition 2); after conditioning on the set of remaining production sites, the entry game with private cost draws becomes the same as in Doraszelski and Satterthwaite (2010) and leads to a similar optimal cutoff rule given by (12), and the dynamic trading decision in (11) is isomorphic to a continuous investment choice with evolving support. On uniqueness, while markets will eventually contain multiple firms, only one potential entrant arrives in each period—ruling out the equilibrium multiplicity that commonly arises in static entry models where several firms simultaneously decide whether to enter—but we cannot rule out that more elaborate dynamic trading strategies not explored here might give rise to multiplicity, so we will assume that the data are generated by a unique equilibrium.

III. Empirical Strategy and Estimation

The empirical strategy to identify and estimate the model of Section II involves three parts. First, we identify demand for offsets from observed land development and transaction prices and quantities over time, using price instruments constructed from cost shifters of offset supply (Section IIIA). Second, to identify supply of offsets, we use maps of observed entry and the environmental characteristics of a market's remaining available land suitable for wetland banking. We correct for selection into wetland banking by forward-simulating value functions as in Hotz et al. (1994); Bajari, Benkard, and Levin (2007); and Pakes, Ostrovsky, and Berry (2007) for both incumbents and potential entrants, to recover the distribution of fixed costs consistent with optimal entry (Section IIIB). Third, we identify the local flood externalities of different wetlands using historical changes in wetland extent and realized flood insurance claims (Section IIIC).

A. Demand for Offsets

We first describe how we obtain local demand for development on wetlands given offset prices. For tractability and to allow for spatial correlation across pixels, we partition pixels i into local watersheds h . For each watershed, we obtain local offset demand, which we will aggregate to market-level demand via (4), using water district acre-to-offset ratios, \tilde{v}_h , to convert developed wetland acres into offsets. Our data allow us to construct pixel transitions over five-year intervals, so we estimate demand at the watershed-by-period level, with periods t given by 1996–2001, 2001–2006, 2006–2011, 2011–2016. We calculate the share of development on private wetlands, $\omega_{ht} = Q_{ht}/W_{ht}$, by dividing the area Q_{ht} of private wetlands in watershed h developed in period t by the total area of private wetlands $W_{ht} = \int_{i \in h} w_{it} w_{i0}$ at the start of t . Taking this observed share ω_{ht} as the conditional probability that $i \in h$ develops a private wetland, we obtain the logit equation (14):

$$(14) \quad \ln \omega_{ht} - \ln(1 - \omega_{ht}) = \theta' X_{ht} + \theta_P \tilde{v}_h P_{ht} + \xi_{ht}$$

for each watershed h and period t , where development choices depend on the average offset price, P_{ht} , and other observable determinants of demand X_{ht} , including period and water district fixed effects, flood zone designations, new development on nonwetlands, lagged development density, and lagged demographics such as median income and population.

Identifying Offset Demand.—As wetland offset prices are partly determined by incumbents' trading decisions, and therefore incumbents' beliefs about unobserved demand shifters ξ_t , equation (14) cannot be estimated without an instrument for price. We consider three sets of instruments for local prices, each based on various cost shifters for offset production.

First, we calculate the average production capacity of historical entrants whose service areas contain h . Intuitively, all else equal, greater sunk capacity due to historical entry should shift market prices downward, acting as a downward cost shifter. While realized capacities are endogenous, these capacities are fixed upon entry and cannot be subsequently adjusted, so when we control for the information known by those entrants, they become excluded shifters of future costs (Berry and Compiani 2023). Because banks produce offsets slowly (over an average of 18 years), our sunk capacity instrument can remain relevant over long horizons. The primary concern is that entrants rely on private information about future unobserved demand shifters; our conversations with bankers indicate that they primarily use forecasts based on public information, such as home prices and historical offset prices.

Second, we build Hausman (1996) instruments from endogenous outcomes in nearby markets as proxies for cost shifters in the market of interest. We use average prices and historical entrant capacity from banks in the same water district but different markets.

Third, we use variation in other public wetland and conservation land, which act as natural cost shifters for offset supply. This creates ideal variation in costs for wetland banks, which vary with available private land and its connectivity to

TABLE 2—ESTIMATED DEMAND FOR DEVELOPMENT ON WETLANDS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Credit price coefficient ^a (θ_p)	-0.34 (0.14)	-1.29 (0.28)	-0.98 (0.26)	-1.10 (0.38)	-1.45 (0.60)	-2.32 (0.58)	-1.06 (0.39)
Implied parameters							
Average price elasticity	-0.3	-1.13	-0.85	-0.96	-1.31	-2.03	-0.96
Std dev price elasticity	0.15	0.58	0.44	0.49	0.66	1.03	0.48
Aggregate consumer surplus (bn US\$)	4.11	1.12	1.67	2.62	2.34	2.37	2.64
Instruments							
Historical sunk capacity		✓	✓	✓			✓
Hausman cost shifters					✓		✓
Government conservation land purchases						✓	✓
Additional controls							
Lagged demographics ^d			✓	✓	✓	✓	✓
HUC8 fixed effects ^e				✓	✓	✓	✓
First-stage <i>F</i> -stat		115.8	117.3	49.8	8.3	21.3	14.3
Observations	758	758	758	758	629	758	629
Adjusted <i>R</i> ²	0.70	0.68	0.70	0.71	0.68	0.64	0.70

Notes: Instrumental variable estimates of (14) at the watershed-by-period level for watershed-periods with prices and nonzero development. See Section IIIA for details. All columns include period and water district fixed effects and controls for baseline flood risk^b and lagged development density.^c Watersheds correspond to HUC12 units. Periods are (1996–2001, 2001–2006, 2006–2011, 2011–2016).

^aPrice coefficient from (14). Scaled by 1/100,000.

^bPercent areas designated storm surge or 100-year flood zones.

^cShare developed and share highly developed.

^dPopulation and median income.

^eHydrological unit code (USGS 2013). All omitted coefficients reported in Supplemental Appendix Table A5.

existing conservation land. Specifically, for each period and watershed, we construct the total area of public wetlands of all other watersheds in the same market (excluding land used by wetland banks). Most of this variation is cross-sectional, though some evolves over time through new land purchased under Florida's conservation buyback programs, Florida Forever and Preservation 2000.

Estimates.—Table 2 reports the demand estimates. The key object of interest is the elasticity of local wetland development with respect to the average offset price. As described above, our empirical strategy instruments for the current offset price using various offset production cost shifters.¹² These instruments vary in strength, with own historical capacity as the strongest instrument, with a first-stage *F*-statistic ranging from 49.8 to 117.3, even conditional on our diverse controls, though the Hausman and public conservation land instruments also meaningfully shift prices (8.3 and 21.3). In addition, these instruments shift prices in the way theory predicts: Markets with larger historical entrants, more historical entrants in neighboring markets, or greater public conservation land, each have lower prices.

¹²Without the instrument, column 1 of Table 2 shows that OLS implies an average price elasticity of demand about -0.3. This is particularly concerning for monopoly and duopoly markets, where incumbents should prefer to locate on a less inelastic part of the demand curve. Several possible sources of upward bias for the OLS coefficient arise in our context. For example, places with greater unobserved values for development may have higher costs of wetland banking.

Columns 2–7 report instrumental variable estimates of (14). Across various controls, the estimated elasticity is close to -1 , showing both a significant relationship between the cost of purchasing an offset and development on local wetlands and that demand is moderately elastic. These findings suggest that these markets are empirically meaningful determinants of land-use decisions. To our knowledge, this is the first estimate of this demand curve, so there is no prior literature for us to benchmark our estimates. We take the estimate in column 3, where $\hat{\theta}_p = -0.98$, as our preferred estimate for subsequent analysis.

B. Restoration Costs

The main identification challenge to recovering unobserved production costs is that banks may enter more often in some markets because their costs in those markets are especially low, because entry in those markets is unusually profitable, or both. Our estimates of offset demand, combined with structure on the entry and trading games, allow us to identify fixed costs using the equilibrium conditions derived in Section IID. In the Markov perfect equilibrium, trading and entry are given by the functions \mathcal{Q} in (11) and ϕ_t in (13). We take the two-step approach of Bajari, Benkard, and Levin (2007). First, we estimate flexible entry and trading strategies as well as production functions for wetland banks. Second, we calculate implied flow payoffs and value functions for incumbents, which identifies the distribution of fixed costs: Conditional on those payoffs, remaining variation in observed entry reflects fixed costs.

Entry.—Our model specifies a finite set of production locations within each market, over which entry opportunities arise at random. Entry therefore depends on sufficient statistics for the remaining production sites, \mathcal{F}_t^c , and the other market conditions contained in s_t . Our data include the location and date of every wetland bank as well as land ownership and characteristics everywhere within each market. To estimate the entry model of (13), we proxy for remaining land available for wetland restoration, \mathcal{F}_t^c , with the areas of public and private wetland and the number of incumbent firms. We estimate annual market-level entry probabilities at the market-year level with the following probit specification,

$$(15) \quad \mathbb{P}(\text{enter}_t | \hat{s}_t) = \Phi(g(\hat{s}_t)),$$

where Φ is the Gaussian CDF and g is a log-linear function in an approximate market state $\hat{s}_t \equiv \{t, \sum_h W_{ht}, \sum_h W_{ht}^{\text{public}}, \sum_h X_{ht}, |\mathcal{F}_t|, \sum_f B_{ft}\}$, defined to include the year; the market's total private and public wetlands, $\sum_h W_{ht}$ and $\sum_h W_{ht}^{\text{public}}$; total population, median income, average flood risk, and water management district, $\sum_h X_{ht}$; the number of incumbents, $|\mathcal{F}_t|$; and total incumbents' reserves, $\sum_f B_{ft}$.

Our probit estimates of (15) indicate that across market-years, entry occurs more frequently in markets with more wetlands, more developed land, more development occurring on nonwetlands, and fewer incumbents. Although χ^2 -tests clearly distinguish our entry model from one that does not vary across markets (Supplemental Appendix Table A6), our probit estimates do not explain all of the observed variation

in entry decisions, so our distributional assumption also acts as an important source of identifying variation.

Production Function.—We observe the numbers and dates of offsets, b_{ft} , issued to each bank directly from various regulatory records. We also observe each bank’s total offset allowance for the life-span of the project, $\tilde{v}_f = \sum_{t \geq 0} b_{ft}$. This is useful for us because the typical bank in our data has not yet produced all of its offsets, given the lags in production. Together with the entry date of the bank, this allows us to construct production as a function of the bank’s age, size, and local characteristics.

We specify the empirical analog of the production function (6) in two pieces. First, we are interested in lifetime production, \tilde{v}_f . Second, we are interested in the timing of offset releases, given by

$$b_{ft} = \mathcal{B}(T_{ft}, \tilde{v}_f) = \sum_{\tau \geq 1} \mathbf{1}_{\{T_{ft}=\tau\}} \alpha_\tau \tilde{v}_f.$$

Our simulations do not estimate \tilde{v}_f or α_τ ; instead, they obtain \tilde{v}_f by drawing from the empirical distribution of $\{\tilde{v}_f\}$ over entrants in the data, then set $\alpha_\tau = 1/10 \cdot \mathbf{1}(\{\tau \leq 10\})$ to approximate the time-to-build discussed in Section ID.

Trading.—We estimate the dynamic trading strategy (11) by predicting trades as a function of a bank’s current reserves and future production, its rivals’ characteristics, and its market’s state. In the data, we observe b_{ft} and q_{ft} , the number of offsets issued to, and sold by, each bank. This lets us estimate trading strategies at the incumbent-year level from 1995 to 2018, via

$$(16) \quad q_{ft} = \chi(\hat{s}_t, B_{ft}, T_{ft}, \tilde{v}_f),$$

where χ is an empirical approximation to (11) and \hat{s}_t is the approximate market state defined after equation (15). As in many applications of Bajari, Benkard, and Levin (2007), the ideal policy function χ consistent with the model is a nonparametric function of a high-dimensional state space, but its estimation in a finite sample may lead to error. In taking (16) to the data, we take three steps to mitigate this concern. First, the rules that allow banks only to trade their certified offsets lead us to bound $\chi_f \in [0, B_{ft}]$. Second, we use a logistic regression to predict a bank’s observed trades as a share of its lifetime production \tilde{v}_f , given the bank’s reserves B_{ft} and age T_{ft} , and its rivals’ characteristics and demand shifters in the approximate market state \hat{s}_t . Third, we discipline the fitted values $\hat{\chi}_f$ from this regression with individual rationality (IR) constraints, $\hat{\chi}_f \leq q_f^{IR}(\hat{s}_t)$, derived from static Cournot first-order conditions using the aggregate demand elasticities, $\eta(Q, \hat{s}_t) \equiv \frac{\partial P(Q, \hat{s}_t)}{\partial Q} \frac{Q}{P(Q, \hat{s}_t)}$, and the vector of equilibrium market shares and credit balances:

$$(17) \quad 1 + \frac{q_f^{IR}}{q_f^{IR} + \sum q_{-ft}} \cdot \eta(q_f^{IR} + \sum q_{-ft}, \hat{s}_t) = 0.$$

These constraints rule out predicted trades that exceed the point at which a bank’s marginal trade would shift the inverse aggregate demand curve downward by so

much as to lower that bank's total profit (i.e., violate static individual rationality). We plot actual trades, predicted trades, and static Cournot constraints, q_{ft} , $\hat{\chi}_f(\hat{s}_t)$, and q_{ft}^{FR} , in Supplemental Appendix Figure A5.

State Transitions.—We model the state transitions of the exogenous demand shifters (local development on other land and lagged demographics) as AR(1) processes. Development on nonwetlands depends significantly on the previous stock of developed land, and population and income are highly persistent. The transitions of the remaining endogenous states—in particular, the extent of private wetland and developed land—are then calculated from entry, production, sales, and these shifters.

Value Functions.—Next, we combine our estimates for entry, trading, and production with our earlier estimates of the regulator's determination of environmental quality and aggregate local demand for offsets to obtain the expected value function via forward simulation. Specifically, given a conditional distribution $H(s_{t+1}|s_t)$ for the transition from state s_{t+1} to s_t , we can calculate the expected value function in (10) as

$$(18) \quad V(B_{f0}, T_{f0}, s_0, \tilde{v}_f) = \sum_{t=0}^T \delta^t \int_{S'} \Pi(\mathcal{Q}(s_t, B_{ft}, T_{ft}, \tilde{v}_f), s_t) dH^t(s_t | s_0),$$

where $H^t(\cdot | s)$ denotes iteration, e.g., $H^2(\cdot | s) = H(\cdot | H(\cdot | s))$, etc., and $T \gg 0$.

We obtain H as the empirical distribution of a large number of sample paths constructed by drawing entrants probabilistically at each t . To then estimate costs, we invert $\phi_t(s_t, x_{ft}) = G_t(z | x_{ft})$ at $z = V(0, 1, s_t, \tilde{v}_f)$ to obtain the conditional entry cost distribution $G_t(\cdot | \tilde{v}_f, \mathcal{F}_t^c)$. Supplemental Appendix C describes the algorithm in detail.

Estimates.—Table 3 reports results for our entry cost estimator. Conditional on entry, we estimate average entry fixed costs of \$8.4 million per bank, or \$29,000 per offset certified (median \$13,700), with considerable dispersion across banks, with an interquartile range of \$5,800 to \$35,500/offset. Notably, these estimates resemble observed costs discussed in Section ID but not used in estimation. Table 3 shows that average observed entry costs (land costs plus restoration costs) obtained from wetland bank contracts are \$5.3 million or \$24,000 per offset (median \$16,000). We take these resemblances to suggest our dynamic cost estimates seem reasonable, given that the two major costs of wetland banking other than unobserved permitting costs should be restoration and the opportunity cost of land.

The structural parameters in Table 3 also provide some additional insight into entry costs. First, unconditional means are much higher than average realized costs, reflecting the fact that entry occurs infrequently. This highlights the importance of correcting for selection into wetland banking. Second, the estimated markups and rates of return on capital appear plausible, averaging 6.1 percent, with an interquartile range of 1.8–8.3 percent, comparable to the average real rate of return of 5.86 percent on US housing from 1980 to 2015 (Jordà et al. 2019, Table 7).

TABLE 3—ESTIMATED WETLAND BANK COSTS

	Observations	Mean	SD	q25	q50	q75
First-stage entry probabilities, $p_{\{enter\}}$	106	0.12	0.07	0.08	0.13	0.16
Value functions, $\mathbb{E}[V]$	106	17.88	29.99	1.37	4.82	15.91
Parameter estimates						
$\mu_{\kappa}(s_{mt})$	106	15.98	2.56	14.46	15.12	16.12
$\sigma_{\kappa}(s_{mt})$	106	1.17	1.68	0.12	0.94	1.19
Implied costs						
Realized entry cost estimate (MM/bank)	106	7.01	10.71	1.06	2.85	6.85
Est. entry costs per credit ('000/bank)	106	29.93	41.97	5.79	13.66	35.53
Implied rate of return on capital (pct)	106	6.08	5.97	1.79	3.81	8.30
Comparison with contract data						
Observed entry costs (MM/bank)	79	5.29	6.09	1.42	2.86	7.18
Observed entry costs per credit ('000/bank/credit)	79	23.95	23.27	9.20	15.99	31.17
Observed construction costs (MM/bank)	86	1.61	2.50	0.36	0.97	1.81
Observed land costs (MM/bank)	95	5.05	10.53	0.57	1.89	5.53

Notes: Wetland bank cost estimates. See Section IIIB for details. Additional results appear in Supplemental Appendix Table A8.

C. Wetlands and Flood Protection

The last aspect of our empirical analysis involves data on environmental outcomes, where we focus on unpriced local flood protection benefits from wetlands. The causal relationship we seek to recover is how—all else equal—altering wetland conservation and restoration will affect the economic costs of flooding in surrounding areas. The ideal research design is to randomly assign wetlands to locations and evaluate flood damages across locations that differ only by their assigned wetlands. However, as we emphasized in discussing the regulations and incentives for land use, wetlands are not randomly developed. The primary threats to identification are unobserved changes that (i) heighten exposure to flood risk and (ii) correlate with changes in wetland extent. We therefore control for each watershed's historical flood claims, prior developed area, and inherent flood risk measured by flood hazard maps. In addition, we observe the source of new development using state transitions for each pixel, which allow us to control for new development on (nonwetland) vacant land, a proxy for unobserved shocks to development payoffs that correlate with both wetland destruction and changing flood risk exposure. Finally, as an outcome, we use only flood claims for structures built prior to 1995, to ensure that our measure primarily reflects the spillovers from wetland protection, not new properties built on wetlands that are (mechanically) exposed to floods.

We assume that flood damage amounts arise according to a conditional Poisson process D_h , where wetlands provide flood protection in proportion to the underlying risk of the local watershed. We opt for a Poisson specification of the conditional mean for three primary reasons. First, flood damages are always nonnegative and often zero; the Poisson distribution has long been viewed as the canonical model for the arrival process of a count of events even when none occur (von Bortkiewicz 1898; Pynchon 1973; Hausman, Hall, and Griliches 1984). Second, if the mean is correctly specified, a Poisson quasi-maximum likelihood estimator (qMLE) for the

conditional average treatment effects does not restrict any other moments, making it fully robust to distributional misspecification (Wooldridge 1999). Third, an exponential functional form for the conditional mean is particularly important for us given that observed flood claims range over eight orders of magnitude across watersheds (Table 1).

We also let wetlands converted to development and wetlands restored through banks differentially affect outcomes because we do not want to assume that these two activities have symmetric effects on flooding; development often replaces wetlands with impervious surfaces, while restoration can improve the functionality of degraded wetlands. Marginal development on wetlands, Q_{ht} , affects expected flooding through a coefficient ζ_d and underlying risk \mathcal{D} ,

$$(19) \quad \frac{\partial}{\partial Q_{ht}} \mathbb{E}[D_h | X_{ht}, Q_{ht}, B_{ht}] = \zeta_d Q_{ht}^{-1} \mathcal{D}(X_{ht}, Q_{ht}, B_{ht}),$$

whereas wetland restoration through banks, B_{ht} , involves marginal changes of $\frac{\partial}{\partial B_{ht}} \mathbb{E}[D_h | \cdot] = \zeta_b \frac{1}{\sqrt{1 + B_{ht}^2}} \mathcal{D}(X_{ht}, Q_{ht}, B_{ht})$. Our baseline specification to estimate $\mathbb{E}[D_h | \cdot]$, ζ_d , and ζ_b uses ex post outcomes across watersheds h —average annual flood damage in the post-period (2016–2020) to structures built prior to 1995—to study wetland changes due to offsets from 1996 to 2016. We use flooding in the pre-period (1991–1995) to control for unobserved confounders. The qMLE Poisson estimator assumes

$$(20) \quad \mathbb{E}[\text{claims}_{h,\text{post}}] = \exp(\zeta_d \ln Q_{h,1996-2016} + \zeta_b \cdot \text{asinh} B_{h,1996-2016} + \varphi(\text{claims}_{h,\text{pre}}) + \gamma' X_h),$$

where $\text{claims}_{h,\text{post}}$ and $\text{claims}_{h,\text{pre}}$ are measured in 2020 dollars; $\zeta = (\zeta_d, \zeta_b)$ are coefficients on wetland development $Q_{h,1996-2016}$ and restoration $B_{h,1996-2016}$; $\text{asinh}(\cdot)$ denotes the inverse hyperbolic sine, $\varphi(x) \equiv \rho_0 \mathbf{1}_{x>0} + \rho_1 x + \rho_2 x^2$; and X_h includes new development on nonwetlands, percent area in baseline flood risk categories (A and V zones), baseline development and high-intensity development densities, and water management district fixed effects.

At least four aspects of (20) are worth noting. First, we follow prior literature to assume that lost wetlands affect floods through their extent or acreage (Brody, Highfield, and Blessing 2015; Sun and Carson 2020). We experimented with some specifications involving additional measures of wetland fragmentation, cluster size, and quality but were unable to detect effects. Second, the constant ζ in (20) implies that level differences in expected local protection arise through the intercept, via differences in historical exposure, $\varphi(\text{claims}_{h,\text{pre}})$, and other local conditions, such as development density and baseline flood risk in $\gamma' X_h$. Third, estimating (20) at the watershed level captures within-watershed externalities of development in h but rules out spillovers to watersheds $h' \neq h$. We test for such spillovers by evaluating the effect of wetland development on flooding in upstream or downstream watersheds; they do not appear empirically relevant here, indicating that the local watershed is an appropriate spatial unit of analysis for our study. Fourth, floods

TABLE 4—ESTIMATED LOCAL FLOOD DAMAGE FUNCTIONS

	(1)	(2)	(3)	(4)
Development on wetlands (ζ_d)	0.428 (0.077)	0.245 (0.083)	0.243 (0.084)	0.147 (0.085)
Wetland bank area (ζ_b)	-0.083 (0.042)	-0.093 (0.036)	-0.093 (0.036)	-0.108 (0.034)
Flood Zone V (storm surge) (%)		2.922 (0.892)	2.919 (0.895)	1.776 (0.949)
Flood Zone A (100-yr) (%)		0.849 (0.824)	0.859 (0.831)	1.390 (0.583)
Nonzero baseline flood claims (1991–1995)		3.069 (0.417)	2.700 (0.395)	2.407 (0.288)
Baseline flood claims (1991–1995)		0.236 (0.097)	0.236 (0.097)	0.081 (0.117)
Baseline flood claims (1991–1995) squared		-0.009 (0.007)	-0.009 (0.007)	0.001 (0.008)
Additional controls				
Demographic controls		✓	✓	✓
HUC8 FEs				✓
Implied damages (\$/acre)				
0%	135.7	6.7	9.7	0.000
10%	5,625.1	873.6	1,065.1	124.2
25%	16,624.0	2,977.9	3,533.3	589.1
50%	54,315.2	10,145.9	12,746.6	2,614.4
75%	184,120.7	39,954.6	52,343.7	11,588.8
90%	391,284.0	151,036.6	184,177.3	38,930.7
95%	574,790.8	275,008.5	369,786.0	89,572.5
97.5%	744,815.1	501,671.4	576,853.7	195,402.1
99%	797,514.3	754,819.5	855,061.2	336,950.8
99.9%	1,076,979.0	4,792,847.0	5,256,975.0	1,531,050.0
100%	10,579,423.0	5,380,962.0	5,369,403.0	2,257,844.0
Observations	1,226	1,226	1,015	1,226

Notes: Quasi-Poisson estimates of (20) at the local watershed level for watersheds with nonzero wetland development. All columns include water district fixed effects and controls for baseline development density and other development on nonwetlands. The outcome is flood insurance claims after the market (2016–2020) for properties built prior to the market (1995); see Supplemental Appendix Table A12 for other outcomes. Column 3 restricts the sample to watersheds with nonzero flood insurance policies in 1995. Implied damages report quantiles of watershed-level expected marginal damages (at observed development) per acre wetland developed. All omitted coefficients reported in Supplemental Appendix Table A9. Robust (HC1) standard errors clustered at the HUC12 level in parentheses.

involve economic damage beyond insurance claims. For example, our measure will not account for flooding of uninsured properties, damage to insured properties that exceed policy limits, or the cost of defensive investments to lower flood risk.¹³

Estimates.—Table 4 presents the results of estimates of (20) across different controls and subsamples. Column 1 shows a strong positive correlation between development on wetlands and flood insurance claims, consistent with the prior literature’s findings, as well as omitted variable bias from underlying hydrological factors that

¹³In 2015, 38.7 percent of Florida households (52.5 percent when weighted by median household income) in flood risk zones had flood insurance (FEMA 2018, Tables 2.3, A4, A5, and A6). In our claims data, coverage limits bind for 5.8 percent of total claims (2.9 percent of building claims and 15 percent of content claims).

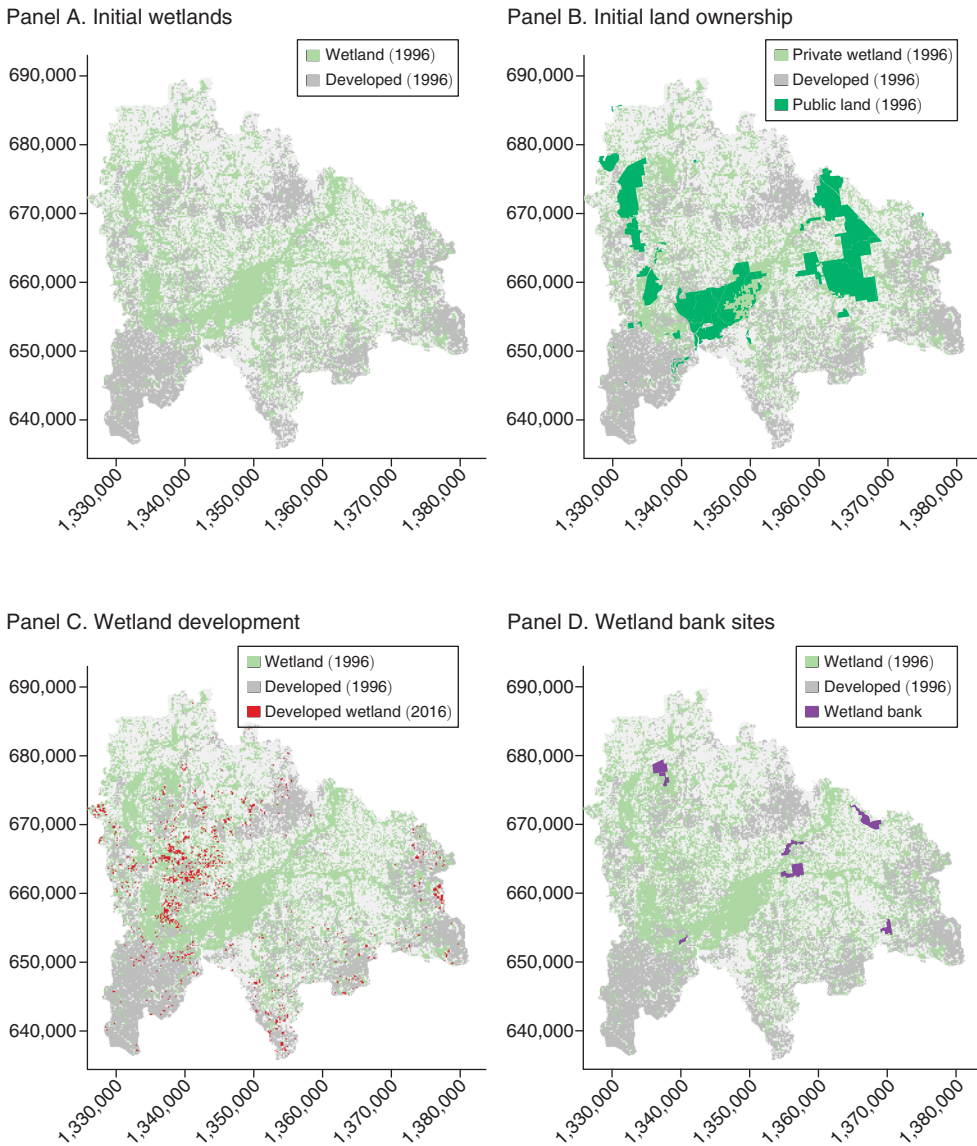


FIGURE 2. INITIAL CONDITIONS, OWNERSHIP, DEVELOPMENT, AND RESTORATION

Notes: An example of our data on land use, ownership, and wetland offsets within a market. Panel A: Initial wetland (green) and developed (gray) pixels in 1996. Panel B: Initial public land (dark green) in 1995. Panel C: New development on wetlands from 1996 to 2016 (red). Panel D: Wetland bank parcels established by 2018 (purple).

make places with more wetland development disproportionately exposed to increasing flood risk. In column 2, which controls for historical flood claims, the coefficient $\hat{\zeta}_d$ falls to 0.245, about half of its value in column 1, and remains significant at the 1 percent level. Flood zone designations also strongly predict damages; the estimates imply that a watershed with 10 percent more of its area in a storm surge flood zone should have 29.2 percent greater expected damages. Column 2 also shows that wetland restoration on bank sites delivers statistically significant flood protection,

with $\hat{\zeta}_b = -0.093$, though, as we show below, this does not translate into especially large protective values because banks locate in relatively few, and not extremely risky, watersheds.

We use (2) as our preferred specification when we evaluate the effects of wetland reallocation on insured flood damages below. Columns 3 and 4 of Table 4 show that similar estimates obtain when we add hydrological region fixed effects or drop watersheds that do not have flood insurance in 1995. The estimates are also robust to using nominal instead of deflated claims, different methods of matching geocoded claims to watersheds, and different windows of average historical flood claims (Supplemental Appendix Table A11). They are also broadly robust to different assumptions about how wetlands affect flooding through the hydrological network.¹⁴ The choice of functional form is important, though not essential; for robustness, we report results that predict realized flood claims directly by transforming flood claims with the inverse hyperbolic sine and controlling for baseline flood claims (Supplemental Appendix Table A10). This alternative specification hews closer to the realized outcomes in the data; it implies similar overall flood damages from offset trade but some interesting distributional differences.¹⁵

Our flood protection estimates compare favorably with some recent work on floods and wetlands, summarized in Supplemental Appendix Table A14. They imply annual flood damage spillovers from development on Florida wetlands averaging about \$1,400/ha, which resemble earlier studies finding average annual wetland flood protection values in the Gulf Coast that translate to \$511/ha (2020 US dollars) in Florida (Brody, Highfield, and Blessing 2015). For high-risk storm flood zone watersheds, we estimate annual flood damages of \$25,200/ha, not dissimilar from recent estimates of \$18,000/ha in storm surge zones (Sun and Carson 2020). A notable outlier is recent work by Taylor and Druckenmiller (2022), whose linear average treatment effects would imply implausibly large increases in flood claims for Florida.¹⁶ The order-of-magnitude discrepancy between our results and theirs likely arise from specification differences. Taylor and Druckenmiller (2022) specify a linear model that they estimate at the zip code level with general-purpose land cover data and no data from flood risk maps. We specify a nonlinear model that we estimate at the watershed level with land cover data designed to study local wetland changes over time and granular maps of flood zone designations. We also take a different approach to measuring spillovers than Taylor and Druckenmiller (2022) (damage to structures built before 1995, not all damages in neighboring zip codes) because it appears to better explain our data (footnote 14).

¹⁴A watershed's location in the hydrological network mildly predicts its flood damages; wetland development in neighboring watersheds does not predict local damages. Supplemental Appendix Table A13 contains the results.

¹⁵First, the distribution of watershed-level marginal flood damages shifts rightward, with a much fatter right tail, reflecting the spikiness of the actual flood damage distribution rather than the conditional mean estimated with Poisson qMLE; second, the estimator cannot detect realized flood protection benefits of wetland banks, in contrast to the small but precise expected benefits estimated with Poisson. Coefficient estimates from a negative binomial model imply a similar relationship between development and floods ($\hat{\zeta}_d = 0.285$) as Poisson, with a somewhat fatter right tail, but much less than the inverse hyperbolic sine.

¹⁶Taylor and Druckenmiller (2022) report annual causal effects of \$12,081/ha of wetlands converted to development and \$8,290/ha of wetland lost in highly developed areas from 2001 to 2016. We calculate in Supplemental Appendix Table A14 that these estimates imply that observed wetland changes over this period should have caused 223 percent and 327 percent of the observed increases in flood claims, respectively.

IV. Evaluating the Market

In this section, we draw together the estimates of local demand, entry costs, and flood protection values to address the key questions posed at the start of the paper. First, we evaluate the market relative to historical conservation rules, in order to assess the private gains from trade (Section IVA) and flood externalities (Section IVB) from the transition to the market-based mechanism. Second, we analyze ways to improve the design of the offset market (Section IVC), given our new estimates of private gains from trade and flood externalities.

A. Gains from Trade

In our model, the private gains from trade equal the difference between private values for development on wetlands and mitigation bank fixed costs, integrated over the range of observed trades. To calculate wetland developer surplus in each local watershed h and period t , we calculate expected consumer surplus by integrating over the logit shocks, which, as in Small and Rosen (1981), has the closed-form solution,

$$(21) \quad \begin{aligned} \hat{U}_{ht} &= \int_{\varepsilon} \max \left\{ u(X_{ht}, \xi_{ht}; \hat{\theta}) - \tilde{v}_h P_{ht} + \varepsilon_1, \varepsilon_0 \right\} dF_{\varepsilon} \\ &= \frac{1}{\tilde{v}_h \hat{\theta}_P} \ln \left(1 + \exp \left\{ \hat{\theta}' X_{ht} - \tilde{v}_h \hat{\theta}_P P_{ht} + \hat{\xi}_{ht} \right\} \right), \end{aligned}$$

which we then aggregate by integrating over the empirical distribution $\{W_{ht}\}$ of privately owned wetlands across watersheds in a regional market,

$$(22) \quad CS_m = \sum_t \sum_{h \in m} \tilde{v}_h W_{ht} \hat{U}_{ht}.$$

Figure 3, panel B plots consumer surplus of each trade in descending order.

To obtain costs of supplying offsets, we calculate realized fixed costs from entrants' conditional cost draws using the value functions and estimated cost parameters,

$$\hat{\kappa}_f = \mathbb{E}[\kappa | \text{entry}_f = 1, x_f] = \frac{1}{G(\hat{V}_f)} \cdot \int_{-\infty}^{\hat{V}_f} k dG(k | x_f),$$

for each bank f , as well as producer surplus, $\sum_{f \in m} (\hat{V}_f - \hat{\kappa}_f)$. Given that entrants do not sell all of their offsets by 2016, we calculate producer surplus, PS_m , with an aggregate marginal producer surplus curve integrated over observed trades as described in Supplemental Appendix C. Figure 4, panel B plots realized producer surplus and costs, ordered by descending producer values.

The realized private gains from trade in market m are then the sum of consumer and producer surplus, given by

$$GFT_m = CS_m + PS_m.$$

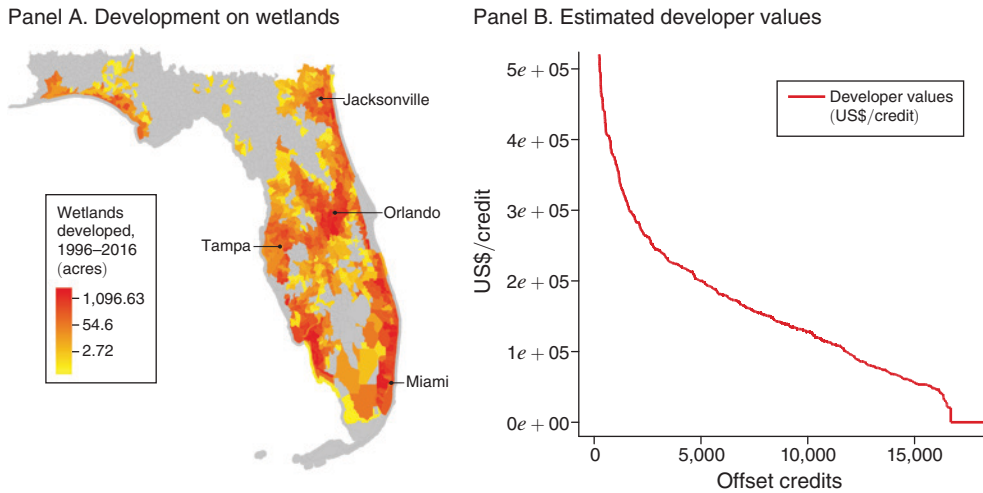


FIGURE 3. DEVELOPMENT ON FLORIDA WETLANDS

Notes: Panel A: Map of local watershed development occurring in offset markets between 1996 and 2016. Local watersheds colored by decile of $\ln(\text{acres of wetlands developed})$. Panel B: Estimated private values for land developers who purchased offsets, calculated with (21), ordered left to right by trades' decreasing estimated value, 1995–2018.

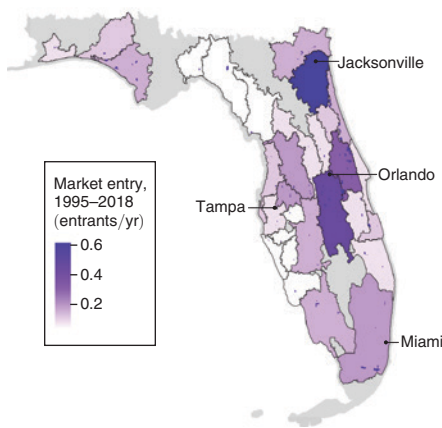
Table 5 reports the results for all of Florida, $\sum_m GFT_m$, our first key empirical finding. The first column shows estimates of developer values, bank costs, and private gains from trade. Developer values, that is, $\sum_{m,t} \hat{U}_{mt}$, equal about \$2.8 billion (2020 US dollars). Total fixed costs, about \$440 million, imply private gains from trade of about \$2.4 billion. Given total sales (\approx \$1.1 billion), consumer surplus from the demand estimates from Section IIIA equals \$1.7 billion, while producer surplus is about \$700 million. These estimates indicate that the private gains from offset trade accrue to both developers of wetlands and wetland banks.

B. Flood Externalities

We now construct marginal environmental externalities using our location-specific estimates of expected flood claims. Given that development of wetlands is irreversible, the social cost of forgone flood protection corresponds an infinite sequence of discounted damages; we scale our annual effect by $\sum_{t=0}^{\infty} 0.95^t$ using a real discount rate of 5 percent in accordance with federal regulatory guidelines during our study period, though we also report totals for 3 percent and 7 percent. We can then obtain marginal damages given by (19) by applying our estimates of $\hat{\zeta}_d$, $\hat{\zeta}_b$, $\hat{\gamma}$, and $\hat{\rho}$ from Table 4 to the data on historical claims and other observables at baseline.

Supplemental Appendix Table A16 reports the distribution of the local flood protection estimates of wetlands across watersheds. The externality from developing a wetland in the first tercile watershed is \$6,600/offset, a rounding error from the viewpoint of a land developer given the typical price of \$88,000/offset. Hence, for many watersheds, wetlands' local flood protection benefits do not justify altering trading rules. However, the highest-percentile externalities (e.g., 90 percent, 95 percentile, of \$792,000/offset and \$1.6 million/offset) exceed observed offset prices.

Panel A. Entry probabilities



Panel B. Estimated producer surplus and costs

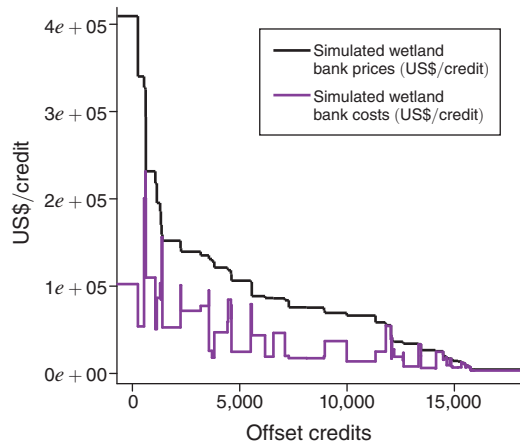


FIGURE 4. WETLAND MITIGATION BANKS

Notes: Panel A: Map of average annual market entry probabilities and wetland bank sites. See Supplemental Appendix Figure A4 for variation across market-years. Panel B: Estimated per credit costs and transaction values for wetland banks, calculated with equations (A2) and (A3) and ordered left to right by increasing simulated price per credit.

This dispersion is also clear from Figure 5, which plots estimated flood damages for each development on wetlands from 1996 to 2016. The jagged blue peaks show high risks in some places amid many wetlands that deliver little or no flood protection value. Figure 6, panel A overlays these estimates with each project's private value. Where the blue spikes cross the red line, development occurred despite flood benefits to conservation that exceed developer values.

Integrating damages over all development and restoration of wetlands, we can approximate the total flood damage from offset trading, which is our second major empirical finding. We find wetlands whose disappearance we attribute to offset trade from 1996 to 2016 would have delivered \$1.7–1.9 billion of flood protection, depending on whether outliers (watersheds above the 99.9 percent, 99 percent, and 97.5 percentile, respectively) are included. Some of these outlier values may reflect measurement or specification error; however, given that the distribution of insured flood damages in the administrative system of record is very fat tailed, it is not unreasonable to expect that the true distribution of marginal local flood protection benefits would also possess a hefty tail. For robustness, Supplemental Appendix Table A11 reports marginal and total damages for some alternative estimates of (20). Both the distribution of flood protection values across wetlands and total flood damages appear similar to the baseline, though the tails above 99 percent appear to be sensitive to the definition of historical flood exposure.

C. Pigouvian Redesign

Finally, we draw on the estimates from our trading and flood protection models to evaluate modified trading rules that account for flood externalities. Our Bajari, Benkard, and Levin (2007) estimator let us recover the industry cost structure, remain

TABLE 5—WELFARE AND OFFSET MARKET DESIGN

	Market	Pigou	Uniform tax
Wetlands developed (acres)	141,606.2	120,097.6	75,168.8
Wetlands offsets used (credits)	16,694.3	14,256.0	8,922.2
Gains from trade			
Developer values (MM)	2,850.6	2,486.2	2,193.0
Supply costs (MM)	440.3	421.4	361.6
Private gains from trade (MM)	2,410.3	2,064.9	1,831.3
Distributional outcomes			
Consumer surplus (MM)	1,678.1	1,339.3	841.6
Producer surplus (MM)	732.2	540.7	124.2
Tax revenue (MM)	0	184.8	865.5
Externalities			
Flood damage (MM)	-1,888.1	-282.1	-714.6
Below 99.9 percentile	-1,888.1	-282.1	
Below 99 percentile	-1,719.5	-282.1	
Below 97.5 percentile	-1,702.4	-284.9	
7% discount rate	-1,132.9		
3% discount rate	-2,643.4		
Welfare (MM)	522.2	1,782.8	1,116.7

Notes: Values in millions of 2020 US dollars. Market outcomes from 1995 to 2020 at the observed equilibrium (column 1), counterfactual equilibrium with local Pigouvian taxes (column 2), and counterfactual equilibrium with the uniform tax that maximizes the sum of private gains from trade and total flood benefits from conservation (column 3). Net present discount values calculated using a 5 percent real discount rate. The uniform tax is calculated to maximize the difference between net surplus and insured flood damages; its optimal level is calculated to be \$97,000/offset.

agnostic toward conduct in the trading game, and evaluate realized trade outcomes. But further counterfactual analysis consistent with our model would require solving for a new dynamic equilibrium, which (i) gives rise to a curse of dimensionality not present when we simulate paths from the observed equilibrium and (ii) requires additional structure to specify the exact form of conduct in the trading game that (11) left indeterminate. Given that solving the dynamic oligopolistic equilibrium model with a large state space is computationally infeasible, our approach is instead to provide reasonable benchmarks for possible counterfactual outcomes, using various approximations that abstract from either dynamics, oligopolistic competition, or full equilibrium adjustments to the higher cost of development under a tax.

Our main counterfactual approximates Pigouvian corrective prices by levying taxes on local wetland developers equal to the expected marginal damage of observed development in their watershed h from (20) and subsidizing wetland bank sales in market m in proportion to average expected flood protection from banks in that market. This reform can be implemented either via a price or a quantity instrument, assessed on either banks or developers.¹⁷

¹⁷Some corrections could also be implemented by altering trading rules. For watersheds with local flood protection values that vastly exceed local developer values, regulators could remove these watersheds from the bank's service area. In addition, in Florida, although state law governs wetland offsets, local governments retain authority to deny permits for wetland development (Grosso and Totoiu 2010).

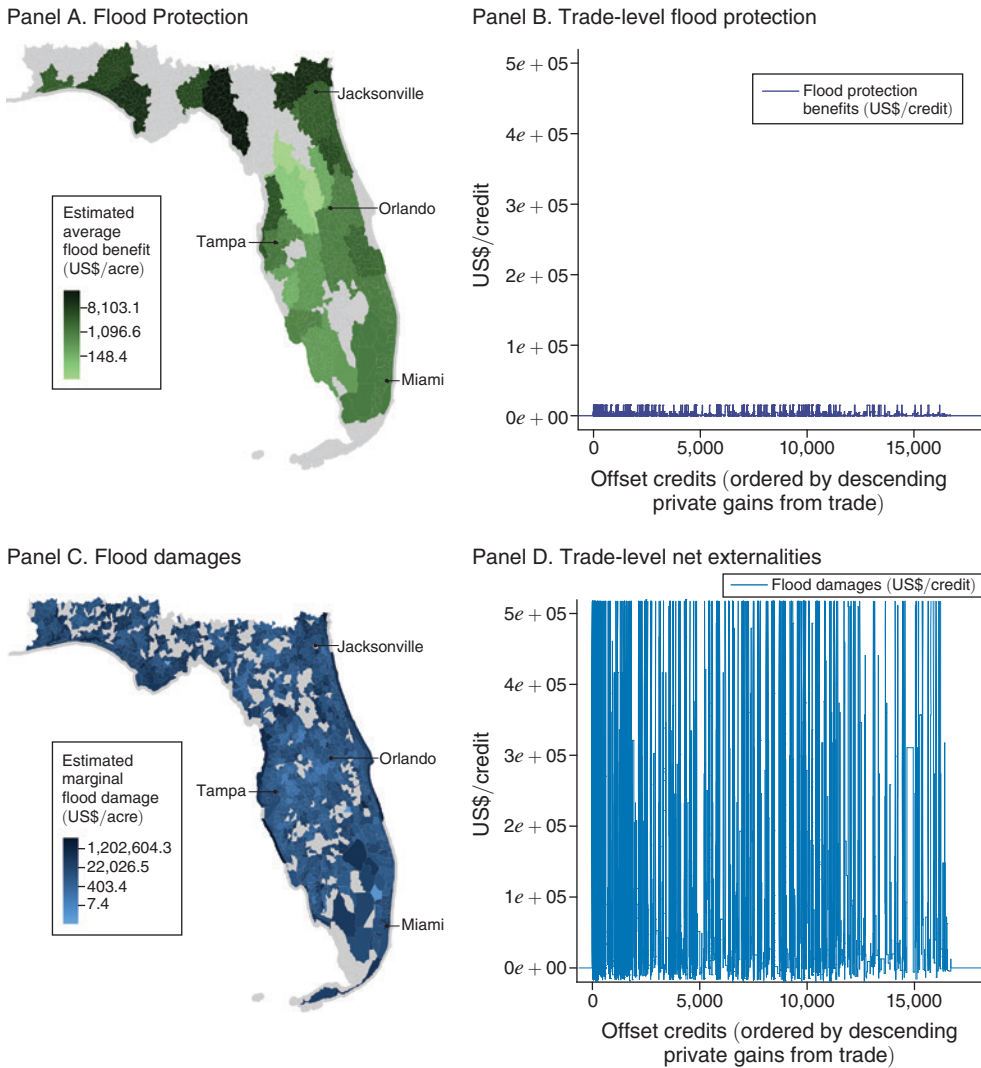


FIGURE 5. REALIZED FLOOD EXTERNALITIES

Notes: Panel A: Map of estimated market-level flood protection benefits from wetland banks. Panel B: Estimated average flood protection benefit for each wetland under the market from 1996 to 2016, calculated with (19) and sorted by descending private gains from trade. Panel C: Map of estimated marginal flood damages at the watershed level for development on wetlands with nonzero wetlands developed. Panel D: Estimated average flood externality for each wetland under the market from 1996 to 2016, calculated with (19) and sorted by descending private gains from trade.

Aggregate demand for offsets, then, reflects the estimated primitives and the counterfactual Pigouvian corrections. On the supply side, we simplify bank behavior in two important and restrictive ways. First, we assume that banks update their trading policy functions with recalculated myopic Cournot constraints (17) at each new state to account for the new aggregate demand they face under the Pigouvian reform, but where these constraints do not bind, they maintain the trading strategy (16) estimated from the observed equilibrium. Relying on the estimated policy functions in the latter case will rule out some important equilibrium responses.

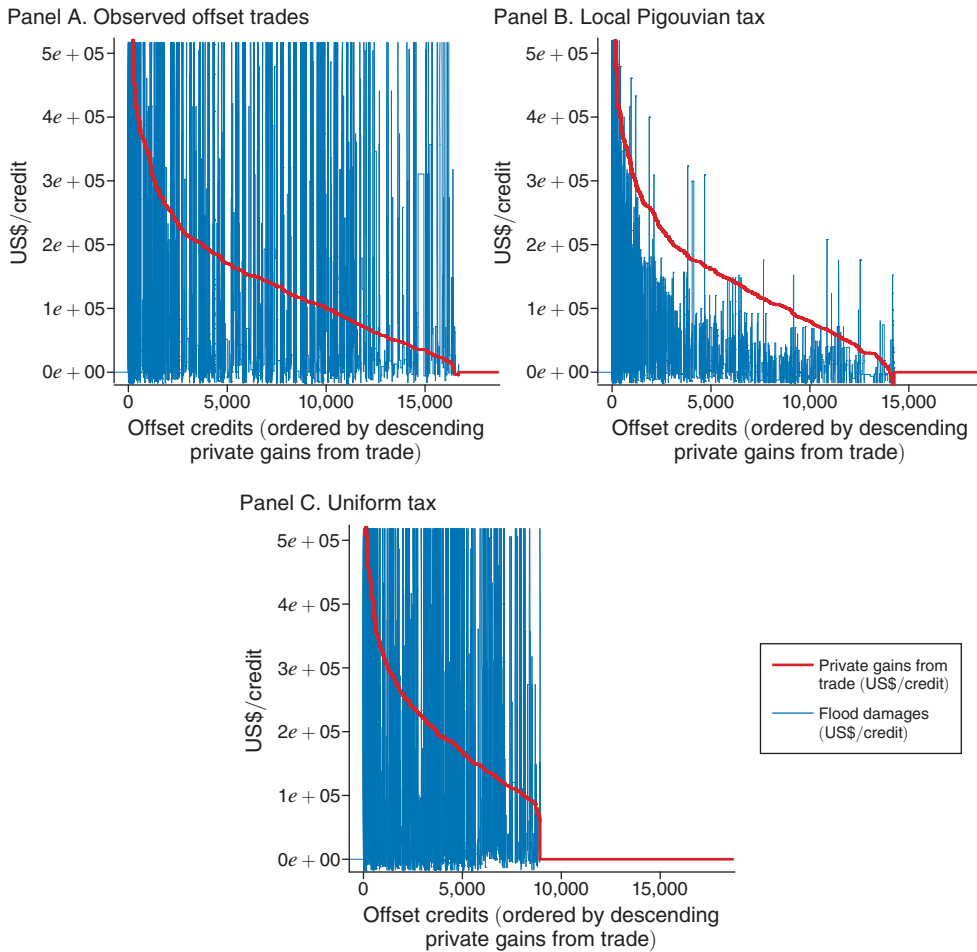


FIGURE 6. PIGOUVIAN REDESIGN

Notes: Panel A: Estimated flood damages from Figure 5, panel D, plotted against the private gains from trade (i.e., the difference between the developer values and bank costs in Supplemental Appendix Figure A7). Panel B: Estimated private gains from trade and flood damages under the Pigouvian flood protection taxes at the local watershed level, sorted by descending private gains from trade. Panel C: Estimated private gains from trade and flood damages under a uniform tax that maximizes the sum of private gains from trade net of total flood damage, sorted by descending private gains from trade. See Section IVC for details.

For instance, firms cannot update their beliefs over future state transitions, such as the evolution of undeveloped wetlands and offset credit balances. Second, we fix the set of incumbent banks to those observed in the data—rather than solving the full dynamic equilibrium—and calculate their producer surplus as in Section IVA. Ruling out lower rates of entry, an obvious extensive margin response to a downward shift in aggregate demand, is also a key limitation that could be relaxed with an exact full-solution method in a simpler model; however, it makes our use of trading policy functions estimated from the observed equilibrium less problematic than if entry also differed in the counterfactual.

The value of this Pigouvian reform given these approximations, reported in the second column of Table 5, is our paper's third major empirical finding. A simple modification of trading rules that accounts for local flood protection benefits—based on observable local characteristics at the USGS (2013) hydrological unit level—lowers excess flood damages by an order of magnitude but preserves more than half of development on wetlands and more than two-thirds of the private gains from trade. Put differently, transitioning to the Pigouvian design creates more than \$4 of flood protection benefits for each dollar of gains from trade forgone ($\frac{1,888 - 282}{2,410 - 2,065} = \frac{1606}{345} = 4.66$). The design also maintains the regulator's No Net Loss goals; the only difference is that it now also accounts for local flood protection.

To isolate the source of the efficiency gains, we also consider an alternative policy that augments the offset market with a uniform flood protection tax on all wetland offset trades in Florida. This policy is of economic interest for at least two reasons. First, comparing the local Pigouvian design with a uniform rule helps to show how heterogeneity in local benefits determines the social value of the reform. For example, if all wetlands delivered the same local flood protection benefits, then the uniform policy should lead to the same trading and flood outcomes as the Pigouvian tax. Second, many environmental policies are constrained to be undifferentiated across place, for various reasons such as simplicity, making it inherently valuable to understand the performance of the second-best corrective policy.

Specifically, we calculate the uniform corrective tax per offset that maximizes total private surplus from trade minus flood damages. The uniform tax that accomplishes this objective turns out to be \$97,000/offset, approximately the mean price in the sample. As the third column of Table 5 shows, such a policy lowers flooding relative to the market but at a much higher private cost. The uniform tax attains about half of the welfare gains from the Pigouvian design but requires a much greater decline in development, lowering development on wetlands to 75,000 acres relative to the 142,000 acres under the market and 120,000 acres under the Pigouvian design. As Figure 6, panel C shows, despite the reduction in development on wetlands, damages exceed private surplus for much of the remaining development, underscoring the need to target local watersheds based on underlying flood risk.

Finally, to analyze the influence of market structure and to assess the sensitivity of our results to the restrictions we imposed on trading responses, Table 6 compares the Pigouvian reform in our baseline model to its performance under three alternative specifications for trading strategies derived under alternative assumptions on conduct in the trading game: (i) full passthrough, (ii) myopic Cournot, and (iii) myopic collusion. As before, we hold fixed the set of banks observed in the data. Assuming full passthrough, reported in Supplemental Appendix Table A17, leads to an upper bound on avoided environmental damages (103 percent of the benchmark's) but a lower bound on the welfare improvement (96 percent of the benchmark's). These differences reflect firms' equilibrium responses through the updated Cournot constraints, which dampen the consumer price shock. Wetland banks prefer to expand supply beyond what would clear the new market at pre-reform prices, rather than to forgo those marginal trades, which lowers the average producer price to 96 percent of its pre-reform level and leads to 7 percent and 5 percent greater consumer surplus and private gains from trade than with complete consumer passthrough.

TABLE 6—WELFARE, PASSTHROUGH, AND OFFSET MARKET STRUCTURE

	Benchmark		Myopic cournot		Myopic collusion	
	Market	Pigou	Market	Pigou	Market	Pigou
Wetlands developed ('000 acres)	141.6	120.1	209.5	176.0	140.6	114.2
Wetlands offsets ('000 credits)	16.7	14.3	24.9	21.2	16.4	13.5
Total transaction volume (MM)	1,172.5	962.1	1,468.6	1,167.7	1,831.9	1,509.0
Passthrough						
Average price ('000\$/credit)	70.2	67.5	58.9	55.2	111.4	111.4
Average price + tax ('000\$/credit)	70.2	80.5	58.9	67.6	111.4	122.8
Producer price change (%)		-3.9		-6.2		-0.02
Consumer passthrough (%)		78.8		70.3		99.8
Gains from trade						
Developer values (MM)	2,850.6	2,486.2	3,969.2	3,395.6	3,470.5	2,885.3
Supply costs (MM)	440.3	421.4	485.1	473.2	425.8	415.8
Private gains from trade (MM)	2,410.3	2,064.9	3,484.1	2,922.4	3,044.7	2,469.5
Distributional outcomes						
Consumer surplus (MM)	1,678.1	1,339.3	2,500.6	1,966.2	1,638.5	1,221.3
Producer surplus (MM)	732.2	540.7	983.5	694.6	1,406.1	1,093.2
Producer surplus (%GFT)	30.4	26.2	28.2	23.8	46.2	44.3
Tax revenue (MM)	0	184.8	0	261.7	0	155.1
Producer surplus change (%)		-26.2		-29.4		-22.3
With lump-sum transfer (%)		-0.9		-2.8		-11.2
Externalities						
Flood damage (MM)	-1,888.1	-282.1	-2,752.0	-554.7	-1,798.0	-253.4
Damages (% pre-reform)		14.9		20.2		14.1
Change (% welfare change)		-465.0		-391.2		-268.6
Change (% GFT change)		127.4		134.3		159.3
Welfare (MM)	522.2	1,782.8	732.1	2,367.7	1,246.7	2,216.1

Notes: Market outcomes from 1995 to 2020 at (i) benchmark offset trading policy functions, (ii) benchmark trading policy functions with local Pigouvian taxes, (iii) myopic Cournot trading policy functions, (iv) myopic Cournot trading policy functions with local Pigouvian taxes, (v) myopic collusion trading policy functions, and (vi) myopic collusion trading policy functions with local Pigouvian taxes. Consumer passthrough (%) is defined as [the average posttax producer price, plus the tax, minus the average pretax producer price] divided by the average tax. Producer surplus change (%) [with lump-sum transfer] defined as the percentage change in producer surplus [plus total tax revenue] relative to previous column. Flooding (% pre-reform) reports the flooding as a percent of flooding in the previous column; changes (%) report counterfactual changes in flood damage relative to the previous column as percentages of the changes in welfare and private gains from trade.

The Pigouvian reform in the myopic Cournot model results in greater declines in equilibrium producer prices, to 94 percent of their average pre-reform level (Table 6, columns 3–4). In contrast, we find nearly complete passthrough when firms collude (Table 6, columns 5–6), consistent with the logit curvature.¹⁸ The reform avoids the most flood damage when passthrough is assumed complete or the trading game is collusive (12.5 percent and 14.1 percent of initial damages, respectively, compared with 14.9 percent and 20.2 percent in the benchmark and Cournot cases). Incidence also differs across trading games. If banks collude, they capture 46 percent of the private gains from trade under the market and 44 percent under Pigou; if banks play their observed strategies with updated χ^{IR} constraints, or play myopic Cournot, they

¹⁸ Even large changes in the small (< 5%) share of wetlands developed will not change the inverse elasticity of demand by much, which means the monopolist's optimal producer price remains nearly constant despite the large shift in aggregate demand under the Pigouvian tax (Supplemental Appendix Figure A8).

capture 30 percent and 28 percent of the private gains under the market and 26 percent and 24 percent under the Pigouvian design.

Our counterfactuals create significant social surplus, which means they may enable varied distributional outcomes depending on how rents are allocated along the transition. For example, redistributing all the tax revenue to producers lump-sum (e.g., by issuing flood protection certificates to firms for free) can almost entirely eliminate the large producer losses under the Pigouvian reform in the benchmark model (moving the decline from 26 percent to only 1 percent), reduce them by an order of magnitude in the Cournot model (from 29 percent to 3 percent), and by half in the collusive case (from 22 percent to 11 percent). Across market structures, the clear winners from the Pigouvian reform are landowners who benefit from the external flood protection. These flood benefits range from 127 percent to 159 percent of the total welfare gains from the reforms and between 2.7 and 4.6 times the forgone private gains from trade, indicating a wide range of transfers from these landowners to existing offset market participants that could make the Pigouvian reform a true Pareto improvement under the assumptions of our model.

D. Broader Lessons and Caveats

We close with some lessons and limitations of our analysis for the design and evaluation of environmental markets. First, our setting features multiple externalities, managed by different government agencies. Many problems involve diverse externalities—for instance, air pollution harms humans and habitats; forests store carbon and kindle wildfires; wolves deliver recreational value but eat livestock. Where market designers lack responsibility over all relevant externalities—for example, when climate change results in unexpected cascades, scientific discoveries reveal previously unknown connections, or new remediation technologies lead to novel externalities—studies like ours seem especially relevant.

Second, we studied regulated offset markets, overseen and enforced by government agencies. Voluntary markets, such as private carbon offset schemes, require other ways to ensure the long-term viability and quality of offsets. A related caveat is that our welfare results assume lawful implementation of (2), which maintains nonflood wetland values. Valuing other wetland amenities (Lupi, Kaplowitz, and Hoehn 2002) lies beyond this study, but if omitted by trading rules, our counterfactuals will affect amenities that differ systematically across wetland banks and developers. For example, if wetland banks deliver fewer ecological benefits than the regulator believes, then our estimated welfare gains will understate the true value of the Pigouvian reform because it leads to greater wetland conservation and fewer wetland banks.

Third, our study highlights some ways in which market design can affect equilibrium offset trade and welfare. The trading zones we study turn out to be wide enough to create flexibility in wetland management (creating private gains from trade), narrow enough (relative to the extent of demand and economies of scale) to make many markets highly concentrated, but (in some cases) too wide to prevent wetlands from relocating to places where they did not deliver the same level of flood protection. Within these trading zones, the extent of competition among suppliers meaningfully affects Pigouvian outcomes.

V. Conclusions

Our paper introduced and applied an empirical framework for evaluating decentralized offset markets. The research design relies on the regulator's certification mechanism, transaction-level market data, equilibrium trading conditions, and auxiliary environmental outcomes. Our approach is applicable to a broad range of environmental markets where the regulator accesses data on offset production (typically required to verify offset quality), the ledger of trades (typically required to avoid double-counting), environmental quality (typically required to enforce environmental laws), and offset prices. We view the framework as particularly useful for analyzing markets for environmental offsets where offset production differs from abatement, where market concentration among offset suppliers seems likely, where verifying offset quality requires long horizons of time, or where concerns exist that some dimensions of environmental outcomes are not fully incorporated into trading rules.

Our empirical findings also have important policy implications. First, regional offset markets created substantial value for participants, despite prohibitions on interregional trade. This economic value primarily arises from the large volume of trade and the high average surplus per trade, reflecting marginal opportunity costs of conservation that considerably exceed new wetland production costs. Second, these offset markets intensified flood damages because wetlands deliver local flood protection benefits that are positively correlated with the marginal opportunity cost of wetland conservation, largely uncorrelated with wetland mitigation banks' incentives to locate, and not included in the current market design. Third, we isolate significant scope for a welfare-improving policy holding fixed the regulator's existing conservation objectives. A Pigouvian tax based on observable local characteristics lowers excess flood damages by more than 80 percent while preserving more than two-thirds of the private gains from trade. Differentiating the market design across watersheds is quantitatively important; a uniform (Florida-wide) tax designed to balance wetlands' flood protection benefits with private gains from trade attains less than half of the benefits of the local Pigouvian design. Market structure also matters, affecting outcomes and incidence of these reforms. We view the robustness of these empirical findings as a key area for future research.

REFERENCES

- Aronoff, Daniel, and Will Rafey. 2026. *Data and Code for: "Conservation Priorities and Environmental Offsets: Markets for Florida Wetlands."* American Economic Association; distributed by Inter-university Consortium for Political and Social Research. <https://doi.org/10.3886/E240508V1>.
- Assunção, Juliano, Robert McMillan, Joshua Murphy, and Eduardo Souza-Rodrigues. 2023. "Optimal Environmental Targeting in the Amazon Rainforest." *Review of Economic Studies* 90 (4): 1608–41.
- Bajari, Patrick, C. Lanier Benkard, and Jonathan Levin. 2007. "Estimating Dynamic Models of Imperfect Competition." *Econometrica* 75 (5): 1331–70.
- Berry, Steven T., and Giovanni Compiani. 2023. "An Instrumental Variable Approach to Dynamic Models." *Review of Economic Studies* 90 (4): 1724–58.
- Borenstein, Severin, James B. Bushnell, and Frank A. Wolak. 2002. "Measuring Market Inefficiencies in California's Restructured Wholesale Electricity Market." *American Economic Review* 92 (5): 1376–405.
- Brevard County. 2018. *Florida Public Record Act: Request for Wetland Offset Purchases, Brevard County Parks & Recreation, Version Reference Number RR 9787*. <https://www.brevardfl.gov> (accessed 10 December 2018).

- Brody, Samuel D., Wesley E. Highfield, and Russell Blessing.** 2015. "An Analysis of the Effects of Land Use and Land Cover on Flood Losses along the Gulf of Mexico Coast from 1999 to 2009." *Journal of the American Water Resources Association* 51 (6): 1556–67.
- Bureau of Economic Analysis (BEA).** 2020. *Annual Gross Domestic Product (GDP) by State: Real Estate and Rental and Leasing. SQGDP2N*. Bureau of Economic Analysis. <https://www.bea.gov/data/gdp/gdp-state>.
- Carlson, Curtis, Dallas Burtaw, Maureen Cropper, and Karen L. Palmer.** 2000. "Sulfur Dioxide Control by Electric Utilities: What Are the Gains from Trade?" *Journal of Political Economy* 108 (6): 1292–326.
- Collier County.** 2018. *Florida Public Record Act: Request for Wetland Offset Purchases*. <https://www.collier.gov/Collier-County/Communications-Government-Public-Affairs/Public-Record-Requests> (accessed 7 December 2018).
- Diamond, Peter A.** 1973. "Consumption Externalities and Imperfect Corrective Pricing." *Bell Journal of Economics* 4 (2): 526–38.
- Doraszelski, Ulrich, and Mark Satterthwaite.** 2010. "Computable Markov-Perfect Industry Dynamics." *RAND Journal of Economics* 41 (2): 215–43.
- Environmental Protection Agency (EPA), and Department of the Army.** 1990. "Memorandum of Agreement between the Environmental Protection Agency and the Department of the Army Concerning the Determination of Mitigation under the Clean Water Act Section 404(b)(1) Guidelines." *Federal Register* 55 (48): 9211–13.
- Ericson, Richardson, and Ariel Pakes.** 1995. "Markov-Perfect Industry Dynamics: A Framework for Empirical Work." *Review of Economic Studies* 62 (1): 53–82.
- Erwin, Kevin L.** 2009. "Wetlands and Global Climate Change: The Role of Wetland Restoration in a Changing World." *Wetlands Ecology and Management* 17 (1): 71–84.
- Federal Emergency Management Agency (FEMA).** 2018. *An Affordability Framework for the National Flood Insurance Program*. Federal Emergency Management Agency.
- Federal Emergency Management Agency (FEMA).** 2020. *National Flood Hazard Layer (NFHL), Version 1.1.1.0*. Federal Emergency Management Agency. <https://msc.fema.gov>.
- Federal Emergency Management Agency (FEMA).** 2023. *National Flood Insurance Program (NFIP) Redacted Claims*. Federal Emergency Management Agency, Federal Insurance and Mitigation Administration. <https://www.fema.gov/about/reports-and-data/openfema/FimaNfipClaims.csv> (accessed 16 January 2023).
- Florida Department of Environmental Protection (FDEP).** 2014a. *FL-SOLARIS/LITS State Owned Lands*. Florida Department of Environmental Protection. https://publicfiles.dep.state.fl.us/otis/gis/data/STATE_OWNED_LAND_SP.zip.
- Florida Department of Environmental Protection (FDEP).** 2014b. *Mitigation Banks*. Florida Department of Environmental Protection. http://publicfiles.dep.state.fl.us/otis/gis/data/MITIGATION_BANKS.zip (accessed July 24, 2020).
- Florida Department of Environmental Protection (FDEP).** 2014c. *Mitigation Banks Service Area*. Florida Department of Environmental Protection. http://publicfiles.dep.state.fl.us/otis/gis/data/MITIGATION_BANK_SERVICE_AREA.zip (accessed July 24, 2020).
- Florida Department of Environmental Protection (FDEP).** 2018. *FL-SOLARIS/CLEAR Conservation Lands*. Florida Department of Environmental Protection. https://publicfiles.dep.state.fl.us/OTIS/GIS/data/CLEAR_CONS_OWNED_LAND_SP.zip (accessed July 12, 2018).
- Florida Department of Environmental Protection (FDEP).** 2022a. *ERP E-Permitting*. Florida Department of Environmental Protection, Division of Water Resource Management, Submerged Lands and Environmental Resources Coordination Program. <https://floridadep.gov/water/submerged-lands-environmental-resources-coordination/content/erp-e-permitting> (accessed March 2022).
- Florida Department of Environmental Protection (FDEP).** 2022b. *Mitigation Bank Ledgers*. Florida Department of Environmental Protection, Division of Water Resource Management, Submerged Lands and Environmental Resources Coordination Program. <https://floridadep.gov/water/submerged-lands-environmental-resources-coordination/content/mitigation-bank-ledgers> (accessed March 2022).
- Florida Department of Transportation (FDOT).** 2018. *Public Record Act Request #P006557-110218*. Florida Department of Transportation. <https://www.fdot.gov/> (accessed 21 November 2018).
- Florida State Legislature.** 2019. § 373.012-373.813. Water Resources.
- Fowlie, Meredith, and Nicholas Muller.** 2019. "Market-Based Emissions Regulation When Damages Vary across Sources: What Are the Gains from Differentiation?" *Journal of the Association of Environmental and Resource Economists* 6 (3): 593–632.
- Fowlie, Meredith, Mar Reguant, and Stephen P. Ryan.** 2016. "Market-Based Emissions Regulation and Industry Dynamics." *Journal of Political Economy* 124 (1): 249–302.

- Gowrisankaran, Gautam, Ashley Langer, and Wendan Zhang.** 2025. "Policy Uncertainty in the Market for Coal Electricity: The Case of Air Toxics Standards." *Journal of Political Economy* 133 (6): 1757–95.
- Grosso, Richard, and Jason Totoiu.** 2010. "Planning and Permitting to Protect Wetlands: The Different Roles and Powers of State and Local Government." *Florida Bar Journal* 84 (4): 39.
- Hausman, Jerry A.** 1996. "Valuation of New Goods under Perfect and Imperfect Competition." In *The Economics of New Goods*, edited by Timothy F. Bresnahan and Robert J. Gordon, 209–48. University of Chicago Press.
- Hausman, Jerry, Bronwyn H. Hall, and Zvi Griliches.** 1984. "Econometric Models for Count Data with an Application to the Patents-R&D Relationship." *Econometrica* 52 (4): 909–38.
- Hotz, V. Joseph, Robert A. Miller, Seth Sanders, and Jeffrey Smith.** 1994. "A Simulation Estimator for Dynamic Models of Discrete Choice." *Review of Economic Studies* 61 (2): 265–89.
- Hsiao, Allan.** 2026. "Coordination and Commitment in International Climate Action: Evidence from Palm Oil." *Econometrica* 94 (1): 1–33.
- Ifrach, Bar, and Gabriel Y. Weintraub.** 2017. "A Framework for Dynamic Oligopoly in Concentrated Industries." *Review of Economic Studies* 84 (3): 1106–50.
- Jordà, Óscar, Katharina Knoll, Dmitry Kushinov, Moritz Schularick, and Alan M. Taylor.** 2019. "The Rate of Return on Everything, 1870–2015." *Quarterly Journal of Economics* 134 (3): 1225–98.
- Keiser, David A., Sheila M. Olmstead, Kevin J. Boyle, Victor B. Flatt, Bonnie L. Keeler, Daniel J. Phaneuf, Joseph S. Shapiro, and Jay P. Schimshack.** 2022. "The Evolution of the 'Waters of the United States' and the Role of Economics." *Review of Environmental Economics and Policy* 16 (1): 146–52.
- Kousky, Carolyn, and Margaret Walls.** 2014. "Floodplain Conservation as a Flood Mitigation Strategy: Examining Costs and Benefits." *Ecological Economics* 104: 119–28.
- Lee County.** 2019. *Florida Public Record Act Request for Wetland Offset Purchases*. <https://www.leegov.com> (accessed January 9, 2019).
- Lupi, Frank, Michael D. Kaplowitz, and John P. Hoehn.** 2002. "The Economic Equivalency of Drained and Restored Wetlands in Michigan." *American Journal of Agricultural Economics* 84 (5): 1355–61.
- Maskin, Eric, and Jean Tirole.** 2001. "Markov Perfect Equilibrium: I. Observable Actions." *Journal of Economic Theory* 100 (2): 191–219.
- McCombs, John W., Nathaniel D. Herold, Shan G. Burkhalter, and Christopher J. Robinson.** 2016. "Accuracy Assessment of NOAA Coastal Change Analysis Program 2006–2010." *Photogrammetric Engineering & Remote Sensing* 82 (9): 711–18.
- Mitigation Marketing.** 2018. *Data on Sales of Wetland Mitigation Credits in the State of Florida*. <https://www.mitigationmarketing.com> (accessed October 24, 2018).
- Muller, Nicholas Z., and Robert Mendelsohn.** 2009. "Efficient Pollution Regulation: Getting the Prices Right." *American Economic Review* 99 (5): 1714–39.
- National Oceanic and Atmospheric Administration (NOAA).** 2020. *NOAA Office for Coastal Management Coastal Change Analysis Program (C-CAP) 1996 Regional Land Cover Data – Coastal United States*. InPort Identifier:gov.noaa.nmfs.inport:48326; Raster Land Cover, 30 m. National Oceanic and Atmospheric Administration. <https://www.fisheries.noaa.gov/inport/item/48326>.
- Pakes, Ariel, Michael Ostrovsky, and Steven Berry.** 2007. "Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples)." *RAND Journal of Economics* 38 (2): 373–99.
- Pynchon, Thomas.** 1973. *Gravity's Rainbow*. Viking Press.
- Rafey, William Minot.** 2020. "Essays in Environmental Markets." PhD diss., Massachusetts Institute of Technology.
- Rafey, Will.** 2023. "Droughts, Deluges, and (River) Diversions: Valuing Market-Based Water Reallocation." *American Economic Review* 113 (2): 430–71.
- Ramsar Convention.** 2018. *Global Wetland Outlook*. Ramsar Convention Secretariat.
- Rapanos versus United States.** 2006. 547 US 715.
- Ruhl, J. B., James Salzman, and Iris Goodman.** 2009. "Implementing the New Ecosystems Services Mandate of the Section 404 Compensatory Mitigation Program." *Stetson Law Review* 38 (2): 251–72.
- Ryan, Stephen P.** 2012. "The Costs of Environmental Regulation in a Concentrated Industry." *Econometrica* 80 (3): 1019–61.
- Sackett versus Environmental Protection Agency.** 2023. 598 US 651.
- Saiz, Albert.** 2010. "The Geographic Determinants of Housing Supply." *Quarterly Journal of Economics* 125 (3): 1253–96.
- Salzman, James, and J. B. Ruhl.** 2006. "'No Net Loss': Instrument Choice in Wetlands Protection." In *Moving to Markets in Environmental Regulation*, edited by Jody Freeman and Charles D. Kolstad, 323–52. Oxford University Press.

- Scott, Paul.** 2013. “Dynamic Discrete Choice Estimation of Agricultural Land Use.” Unpublished.
- Silverstein, Jonathan.** 1994. “Taking Wetlands to the Bank.” *Boston College Environmental Affairs Law Review* 22 (1): 129–61.
- St. Johns River Water Management District (SJRWMD).** 2022a. *SJRWMD ePermit Portal*. St. Johns River Water Management District. <https://permitting.sjrwm.com/ep/> (accessed March 2022).
- St. Johns River Water Management District (SJRWMD).** 2022b. *SJRWMD Mitigation Banking*. St. Johns River Water Management District. <http://webapub.sjrwm.com/agws10/mt/> (accessed March 2022).
- Small, Kenneth A., and Harvey S. Rosen.** 1981. “Applied Welfare Economics with Discrete Choice Models.” *Econometrica* 49 (1): 105–30.
- South Florida Water Management District (SFWMD).** 2016. *Water Management District Boundaries*. South Florida Water Management District. https://services1.arcgis.com/sDAPyc2rGRn7vf9B/arcgis/rest/services/BDWMD_Boundary_Area/FeatureServer.
- South Florida Water Management District (SFWMD).** 2022. *SFWMD RegPermitting*. South Florida Water Management District. <https://www.sfwmd.gov/regpermitting> (accessed March 2022).
- Southwest Florida Water Management District (SWFWMD).** 2022a. *ERP Mitigation Bank Wetland Credit Ledgers*. Southwest Florida Water Management District. <https://www.swfwmd.state.fl.us/media/752> (accessed March 2022).
- Southwest Florida Water Management District (SWFWMD).** 2022b. *SWFWMD ePermitting*. Southwest Florida Water Management District. <https://www.swfwmd.state.fl.us/business/epermitting> (accessed March 2022).
- Souza-Rodrigues, Eduardo.** 2019. “Deforestation in the Amazon: A Unified Framework for Estimation and Policy Analysis.” *Review of Economic Studies* 86 (6): 2713–44.
- Stavins, Robert N., and Adam B. Jaffe.** 1990. “Unintended Impacts of Public Investments on Private Decisions: The Depletion of Forested Wetlands.” *American Economic Review* 80 (3): 337–52.
- Sun, Fanglin, and Richard T. Carson.** 2020. “Coastal Wetlands Reduce Property Damage during Tropical Cyclones.” *Proceedings of the National Academy of Sciences* 117 (11): 5719–25.
- Sunding, David, and David Zilberman.** 2002. “The Economics of Environmental Regulation by Licensing: An Assessment of Recent Changes to the Wetland Permitting Process.” *Natural Resources Journal* 42 (1): 59–90.
- Taylor, Charles A., and Hannah Druckenmiller.** 2022. “Wetlands, Flooding, and the Clean Water Act.” *American Economic Review* 112 (4): 1334–63.
- Turner, Matthew A., Andrew Haughwout, and Wilbert van der Klaauw.** 2014. “Land Use Regulation and Welfare.” *Econometrica* 82 (4): 1341–403.
- US Army Corps of Engineers versus Hawkes Co.** 2016. 578 US 590.
- US Census Bureau.** 2000. *Decennial Census*. US Department Of Commerce, <https://www.census.gov/data/developers/data-sets/decennial-census.2000.html>; <http://api.census.gov/data/2000/dec/> (accessed January 6, 2025).
- US Census Bureau.** 2009–2021. *American Community Survey 5-Year Data*. US Department Of Commerce. <https://www.census.gov/data/developers/data-sets/acs-5year.html>, <http://api.census.gov/data/2009/acs/acs5>, <http://api.census.gov/data/2014/acs/acs5>, <http://api.census.gov/data/2019/acs/acs5> (accessed 6 January 2025).
- US Geological Survey (USGS).** 2013. *Federal Standards and Procedures for the National Watershed Boundary Dataset (WBD)*. 4th ed. US Geological Survey and US Department of Agriculture, Natural Resource Conservation Service.
- US Geological Survey (USGS).** 2022. *Watershed Boundary Dataset (WBD), US Geological Survey, US Department of the Interior, 2-Digit Hydrologic Unit 03, Version 20221206*. https://prd-tnm.s3.amazonaws.com/StagedProducts/Hydrography/WBD/HU2/Shape/WBD_03_HU2_Shape.zip (accessed 13 February 2023).
- von Bortkiewicz, Ladislaus.** 1898. *Das Gesetz der Kleinen Zahlen*. B.G. Teubner.
- Wing, Oliver E. J., Paul D. Bates, Andrew M. Smith, Christopher C. Sampson, Kris A. Johnson, Joseph Fargione, and Philip Morefield.** 2018. “Estimates of Present and Future Flood Risk in the Conterminous United States.” *Environmental Research Letters* 13 (3): 034023.
- Wooldridge, Jeffrey M.** 1999. “Distribution-Free Estimation of Some Nonlinear Panel Data Models.” *Journal of Econometrics* 90 (1): 77–97.
- Zillow.** 2020. *ZHVI All Homes (SFR, Condo/Co-op) Time Series, Smoothed, Seasonally Adjusted*. Zillow Group. https://www.zillow.com/research/data/Zip_zhvi_uc_sfrcondo_tier_0.33_0.67_sm_sa_mon.csv.