

NBER WORKING PAPER SERIES

CONSERVATION PRIORITIES AND ENVIRONMENTAL OFFSETS:
MARKETS FOR FLORIDA WETLANDS

Daniel Aronoff
Will Rafey

Working Paper 31495
<http://www.nber.org/papers/w31495>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2023

The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research. During the period this study was undertaken, members of Daniel Aronoff's immediate family held investments in two wetland mitigation banks in the State of Florida.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Daniel Aronoff and Will Rafey. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Conservation Priorities and Environmental Offsets: Markets for Florida Wetlands
Daniel Aronoff and Will Rafey
NBER Working Paper No. 31495
July 2023
JEL No. L0

ABSTRACT

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 a small number of long-lived producers that restore wetlands over time. We find that offsets led to substantial private gains from trade, creating about \$2.2 billion of net surplus from 1995–2018 relative to a historical conservation mandate. Offset trading also led to large differences in hydrological outcomes, driven by significant differences between restored and existing wetlands in terms of area and location. A locally differentiated Pigouvian tax on offset transactions would have prevented \$1.3 billion of new flood damage while preserving more than two-thirds of the private gains from trade.

Daniel Aronoff
MIT Media Lab
Massachusetts Institute of Technology
75 Amherst St.
Cambridge, MA 02139
daronoff@mit.edu

Will Rafey
Department of Economics
UCLA
8367 Bunche Hall
Los Angeles, CA 90095
and NBER
rafey@ucla.edu

1 Introduction

Environmental offsets 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 the public good. A set of potential producers access restoration opportunities that differ in cost as well as 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 outcomes, which are measured separately. 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 storage and 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 tradeoffs in the United States by mandating “No Net Loss” in existing wetlands. The current rules permit development on local 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 who need to offset development on protected wetlands.

To analyze these markets, we obtain new data on markets for wetland offsets in Florida, where 29% of land by area is wetlands and real estate comprised nearly one-fifth (19%) of the state’s \$1 trillion GDP in 2020 (BEA, 2020). We start by documenting

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

some new stylized facts about wetlands trading. First, we find considerable trading volumes, with more than \$1.2 billion of transactions in regional markets from 1995–2018. 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. For example, we use variation in the issuance of wetland offsets to historical incumbents based on bank-level production schedules fixed at the time of entry, as well as Hausman instruments from other markets in the same region, and variation in the extent of public wetlands that affect the feasibility of offset production.

Second, we estimate a model of industry dynamics of offset supply, using (i) administrative data on the set of operating wetlands 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 *et al.* \(2007\)](#) and [Pakes *et al.* \(2007\)](#) 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 site characteristics. To account for offset storage, we characterize trading decisions as an optimal inventory problem. We then combine estimates of offset demand with optimality conditions for entry, which allow us to obtain the expected profits of an incumbent firm. We then estimate conditional entry cost distributions to rationalize observed entry decisions as solutions to each producer’s dynamic optimization problem as in [Bajari *et al.* \(2007\)](#). In particular, we obtain conditional entry cost distributions that control for local characteristics that affect the feasibility of wetlands restoration across space.²

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

²This use of such resource maps to obtain more realistic approximations of potential entry opportunities echoes recent work using maps of oil reserves in studies of dynamic resource extraction with collusion ([Asker *et al.*, 2019](#)) or strategic learning ([Hodgson, 2018](#)).

Florida, approximately \$700 billion of assets lie in a 100-year flood zone and scientists expect climate change to increase the probability of extreme flooding (Wing *et al.*, 2018). Moreover, new empirical research suggests that the value of these local flood protection benefits may be considerable (Brody *et al.*, 2015; Sun and Carson, 2020; Taylor and Druckenmiller, 2022). 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 development for 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 90% 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 *et al.*, 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 *et al.*, 2016; Fowlie and Muller, 2019; Hernandez-Cortes and Meng, 2020).³ Our findings of the private gains from trade in these regional markets provide substantive evidence of the value that offsets can deliver to landowners and communities, despite rules that prohibit regional trade. Importantly, while these markets can cause other environmental changes and social costs, we show that in our context, addressing the additional externalities can occur without sacrificing most of the private gains from trade.

³Our work also relates to the economics of the private provision of public goods (Bergstrom *et al.*, 1986) and voluntary offsets in particular (Kotchen, 2009), as well as payments for ecosystem services (Jack *et al.*, 2010). Like these papers, we study private incentives to supply and consume offsets; here, however, offsets are mandated and verified by the regulator, not voluntarily purchased.

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.* (2019), as well as recent models of dynamic land use (e.g., Scott, 2013; Hsiao, 2021), typically analyze environmental regulation through price changes, where the only conservation option is to not develop land.⁴ These models rule out the use of land to supply new environmental protection, which is an economically important aspect of contemporary strategies for conservation. Here, we specify and estimate the production technology for new conservation projects, derive equilibrium outcomes within the concentrated markets that arise from the time-to-build and the large fixed costs of these technologies, and endogenize landowners' opportunity costs of meeting a given conservation objective through the offset market.

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 flood risk—not fully incorporated into the original policy design—follows Aronoff and Rafey (2020), who built on recent insights suggesting important interactions between land conservation decisions and flood risk (Kousky and Walls, 2014) as well as recent work connecting land use data with flood outcomes (Brody *et al.*, 2015; Sun and Carson, 2020).⁵ Of particular note is Brody *et al.* (2015), who are the first to construct land cover data that relates 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 *et al.* (2015). Like these papers, our work emphasizes the spillovers created by wetlands that protect existing property,⁶ and our research design relies on detailed hydrological and historical data.

⁴Our work also relates to the economics of land use restrictions (e.g., Saiz, 2010; Turner *et al.*, 2014). The welfare costs to land developers required to conserve wetlands under existing environmental laws act as land use restrictions that we find to be economically meaningful; at the same time, we show that offset markets can help to reduce the costs of environmentally-motivated restrictions on land use without compromising aggregate conservation priorities.

⁵While legal scholars have long noted existing rules for wetland mitigation banking may be imperfect (e.g., Silverstein, 1994), wetland banking has received limited attention from economists (Polasky, 2002), with a few exceptions using surveys (e.g., Lupi *et al.*, 2002; Johnston *et al.*, 2002), simulations (e.g., Fernandez and Karp, 1998), or site evaluations (Boyd and Wainger, 2002).

⁶This textbook externality has a long history; see, e.g., Samuelson (1976) (“Suppose sowing the land to short-lived pine trees prevents floods 500 miles downstream. . . . If . . . foresters and conservationists had brought into court an elaboration of the respects in which forestry is an activity beset with important externalities, carefully and objectively described, Ph.D.’s in economics would be found on both sides of the case. . . . Indeed, if the externalities involved could be shown to be sufficiently important, I am naïve enough to believe that all economists would be found on the side of the angels, sitting thigh next to thigh with the foresters.” (pp. 467–468)).

We build on this prior work 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 effects on flood outcomes, and assess the welfare consequences of incorporating flood externalities into the design of these markets. This paper is the first economic analysis to attain these objectives.

Two, we improve the precision of flood damages models by (a) using a nonlinear model that more closely fits the data on flood damages and (b) focusing on damage to properties built prior to the offsets market to reduce bias. We find that wetlands deliver substantial spillovers in some places, which can affect policy prescriptions. Our findings differ considerably from recent U.S.-wide average 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 supports Taylor and Druckenmiller (2022)’s caveats about heterogeneity across wetland values, and 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 2 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 3 specifies a model of equilibrium supply and demand for wetland offsets and Section 4 describes the empirical strategy and benchmark estimates. Section 5 evaluates private gains from trade, local flood outcomes, and some counterfactual market designs to internalize flood risk; Section 6 concludes.

2 Background and data

2.1 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.⁷ At the

⁷For example, wetlands can purify water, enable recreation activities, and sustain diverse species, services that differ along a range of characteristics, such as location, age, maturity, and salinity.

same time, their conservation precludes alternative land uses and therefore can entail substantial economic cost, often born by landowners whose property includes wetlands.

Activities that risk degrading local wetlands have been regulated at the federal level since the 1973 Clean Water Act. Section 404 of the Clean Water Act prohibits economic activity that risks “significantly degrading” existing wetlands. This prohibition of “significant degradation” has been taken as a mandate to conserve an aggregate stock of ecological and hydrological functions delivered by wetlands. Under this approach, known as the “No Net Loss” principle, wetland degradation can occur legally if it is accompanied by approved actions that “offset” the degradation (Army-EPA, 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 non-market approach was heavily criticized by private landowners and environmental groups alike. Land developers argued that permitting requirements were unduly burdensome (Sunding and Zilberman, 2002), while environmental stakeholders argued that the growing sprawl of on-site mitigation activities did not compensate fully for wetland loss, criticizing regulators as unable to enforce or even verify the success of such mitigation activities (Ruhl *et al.*, 2009).

Tradeable offsets arose in response to these concerns as an innovative way to more flexibly comply with conservation mandates while incurring fewer costs from forgone land development. Rather than requiring land developers whose land included wetlands to undertake on-site mitigation actions or prohibiting development outright, regulators began to allow such landowners to buy offsets from off-site wetland restoration projects, known as “wetland mitigation banks.” These projects commit land to the public trust, and engage in a range of conservation activities to restore degraded wetlands or create new ones (e.g., converting farmland back to its natural state (Erwin, 2009)), often offsetting dozens or even hundreds of new developments.

Our empirical analysis focuses on offsets required by Florida state law. With the greatest percentage 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. Florida regulations governing wetland banking date from February 1994 and encompass all wetlands regulated under the Clean Water Act as well as additional wetlands, such as those not connected to the Atlantic Ocean or the Gulf of Mexico by navigable waters.⁸

⁸The Environmental Reorganization Act of 1993 (Section 373.4135, F.S.) directed the Florida

2.2 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 regulatory approval before either developing protected wetlands or restoring degraded wetlands. To this end, the regulator defines exchange values between restored wetlands and existing wetlands, through on-site assessments and a uniform assessment metric.⁹ 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 1A). These market boundaries, known as wetland mitigation bank service areas, approximate hydrological regions and extend far beyond the local project site.

For the creation of new wetland mitigation banks, the regulator requires an environmental audit, a set of proposed restoration activities, and a detailed implementation schedule. Each project’s total lifetime output reflects the regulator’s assessment of its contribution to “wetland functionality,” which the regulator quantifies on a project-by-project basis. As wetland banks act as substitutes for direct conservation, committing land to a wetland bank requires a permanent conservation easement, ruling out alternative future land uses. In addition, assessment over time creates a delay between entry and obtaining offsets. Total lifetime production is specified at the time of entry, with offsets released gradually as wetlands regenerate over time and 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

Department of Environmental Protection (FDEP) and regional water management districts to adopt rules governing wetland mitigation banking, with FDEP’s rule taking effect on February 2, 1994. The Florida rules largely resemble federal rules.

⁹To define equivalent units across diverse wetlands, regulators use the “uniform mitigation assessment method” (UMAM), which establishes fixed exchange ratios across wetland attributes that deliver a scalar measure of wetland value. The UMAM score captures the “ecological and hydrological functions” the wetland delivers to the surrounding region (Florida State Legislature, 2019, §373.4136(1)). Restoring ecological functions can involve activities such as planting trees, creating habitats, and controlling invasive species. Hydrological functions can entail building dams, bridges, and canals.

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.¹⁰

2.3 Data sources

To study this mechanism for offset production and trade, we develop a new dataset to track wetlands, development, and offsets across Florida from 1995–2018. Our work draws on several new primary sources summarized in Table A1 and detailed in 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. These agencies regulate the creation and sale of environmental 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 1995–2018. In addition, we obtain detailed producer-level data for every wetland mitigation bank operating over this period. 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 at which the initial contract was signed, 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. We supplement the data on these private transactions with Freedom of Information Act requests to county officials and the Florida Department of Transportation for government offset purchases. While transaction prices are not reported to the regulator, our final data includes the majority of trades and nearly the entire period (1998–2018).

¹⁰For 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.

Third, we use maps that track evolving environmental characteristics of coastal land to measure wetland location, extent, and quality. These landcover maps are derived from satellite and aerial data in the National Oceanic and Atmospheric Administration’s (NOAA) Coastal Change Analysis Project (C-CAP) dataset, cover all of Florida at a 30m×30m resolution in 1996, 2001, 2006, 2011, and 2016. This data contains more than 194 million pixels for each of five periods, 136 million of which are contained in watersheds covered by offset trading zones. It gives us an unprecedented view of the evolution of land use in Florida.¹¹

Fourth, we use maps of land ownership. To delineate between private and public land, we use boundaries for all land owned by local, state, and federal entities at baseline (1995) as well as Florida state government conservation land purchases from 1990–2020 under the Preservation 2000 and Florida Forever programs. We intersect ownership with our land cover data to isolate changes in wetlands on private land. We also use annual ZIP-code-level home values from Zillow (1998–2020) and population, income, housing units, and home values from the U.S. Census (2000, 2007–19).

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 1978–2020, which include the claim location, date, and amount, as well as data obtained through a FOIA request that includes total policies held from 1975–2018. In addition, we calculate local measures of inherent flood risk using flood zone designations from the National Flood Hazard Layer (NFHL), which is a product of FEMA. 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 Appendix B. Specifically, we use hydrological boundaries from the United States Geological Service (USGS, 2013) to produce a consistent panel of local watersheds and markets across time that aligns with both hydrological realities and market boundaries (Figure A2). Local watersheds are typ-

¹¹The NOAA C-CAP data offers two major advantages. First, it is specialized to coastal systems in the Gulf Coast: six of its twenty-five land use categories are wetland subtypes. Second, C-CAP is ground-truthed to changes over time, with external validation finding high levels of accuracy (McCombs *et al.*, 2016). Other spatial data, by contrast, update to maximize their most recent edition’s accuracy, which exposes panel analyses to the possibility that they mistake changes in satellite instrumentation, sampling procedures, or classification algorithms for actual changes in land use.

ically about 24,000 acres (40 square miles). Florida contains 1,378 such watersheds, 1,004 of which are contained within offset markets (Figure A1) and included in our final dataset.

2.4 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.

2.4.1 Initial wetland extent and land ownership

We observe the initial condition by constructing exact wetland locations in 1996. Table 1 shows that, for the 136,302,645 pixels in our dataset, 36.4% of them were initially a wetland. When aggregated to local watersheds, the average watershed contained an average of 10,818 acres of wetland (33.2% of its area) at baseline.

Many, but not all of these wetlands will be prospective sites for development, depending on the initial ownership of the land.¹² We delineate public and private land in each watershed by intersecting the hydrological boundaries from USGS with land ownership boundaries from the Florida government. Table 1 shows that much of the land that watersheds include is private: on average, 12.5% of a watershed’s land is conserved as public land in 1995 and the median watershed is 2.1% public. Wetlands are more likely to be publicly owned: 17.5% of a watershed’s wetlands on average, and more than a quarter for the 75th percentile watershed, is public land. At the same time, there are substantial privately-owned wetlands in Florida, with about 7,400 acres of wetland in the average watershed-period.

2.4.2 Spatial patterns in development and wetland restoration

Our data shows significant differences between the places producing offsets and the places buying offsets. To determine where offsets are produced, we match wetland bank entry locations to watersheds. To determine where offsets are bought, we use the conversion of private wetlands into developed land.¹³

¹²Public wetlands—designated as conservation land by local, state, or federal authorities—are not able to be developed by law in most contexts, and therefore must be separated from the analysis of wetland offset demand. Similarly, wetland banks can only enter on private land. At the same time, wetland banks receive additional offsets when they locate near existing public conservation land.

¹³Almost no conversion occurs on public wetlands, with private wetlands accounting for 99.2% of wetland pixels developed from 1996–2016. In our structural analysis, we analyze outcomes only for

Table 1 reports an average of 207.5 wetland acres converted per watershed over the sample period. The median watershed sees very little wetland development (16.4 acres), whereas the 75th percentile watershed sees 186.7 acres converted. Private developers convert an original wetland with probability of 0.037 in the average watershed, with much of this development concentrated in the upper tail: the 75th percentile watershed has a conditional development probability of 0.039, and private developers in the maximal watershed convert more than half of their wetlands (a probability of 0.57).

To illustrate the typical pattern of reallocation, Figure 1 Panel C plots our within-pixel data for a typical market. Wetland development (red pixels) occurs nearby 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 thirty markets that we study, as reported in Figures A4.1–30. Qualitatively, these patterns show that offset trading involves significant spatial reallocation of wetlands within each market.

To quantify these spatial patterns, Table 2 compares watersheds that contain wetland banks with the watersheds that involve the development of substantial areas of initial wetlands, excluding those which contain both. This comparison reveals large differences between where most wetlands are developed and where most wetland banks locate. First, most wetland development occurs in places with greater initial development density: 32.5% of the area of the median high-development watershed starts as developed, vastly exceeding the median watershed’s 4.7% or the median wetland bank watershed’s 3.0%. 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.

Second, wetland mitigation banks enter in watersheds with more initial public wetlands. The average wetland bank watershed contains 13,700 acres of public wetland, compared with the average Florida watershed of 3,300 wetland acres. This pattern is consistent with regulatory incentives that award additional offsets to banks to restore existing wetlands nearby existing conservation land,¹⁴ as well as our evidence from bank project sites that show bank parcels themselves consist of large quantities of initial wetland pixels, and the fact that wetland restoration is much easier adjacent to existing

wetlands converted in places where we observe offsets trade; this corresponds to about two-thirds of all development on wetlands, reflecting the fact that some markets do not begin until after 2000.

¹⁴“Mitigation banks and offsite 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).

wetlands, because it is possible to route water from those wetlands to the newly restored land. The tendency of wetland banks to locate in places that already have large areas of wetland conservation land reflects both regulatory incentives—banks receive additional offset credits for locating near other conservation land—as well as economic incentives, given that land values can be lower in such peripheral places.

Third, in terms of flood insurance, high-development places without wetland banks have greater insured value at baseline (\$18.8 million) than places with only wetland banks (\$10.1 million) or the average watershed in the sample (\$7.2 million). This reflects a positive correlation between more intensive offsets market activity and underlying value at risk. High-development and wetland bank watersheds have similar average historical flood insurance claims (\$413,000/year versus \$314,000/year), but the immense dispersion of these distributions, which have coefficients of variation greater than 5, driven mainly by the extreme values in the right tail, means that none of these average differences are statistically significant.

Summarizing, banks are more likely to choose locations with greater wetland area, lower density of development, and more public wetlands. These correlations are consistent with private landowners searching for low-cost places for restoration, as well as natural constraints on wetland restoration that make some places infeasible for large-scale wetland mitigation bank projects. Developers, in contrast, locate in watersheds with preexisting development and more other development occurring contemporaneously. Both banks and developers appear more frequently in places with greater average annual flood insurance claims and total insured value than the typical watershed.

2.4.3 Offset releases and sales

Banks produce large quantities of offsets relative to the size of their markets and the size of the buyers of offsets. The median bank produces about 210 offsets over its lifetime (or 420 on average), with an interquartile range of 85 to 520 offsets. 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 proportionally with the total area of the wetland bank project site; on average, the ratio of acres to offsets is about 4.5 acres, ranging between 3.6–5.3 across water management districts.¹⁵ Banks also take time to build, reflecting the time required to

¹⁵We also intersect our data on pixel-level initial conditions and our within-pixel transitions with these project sites and find that wetland mitigation banks obtain offsets primarily in proportion to their initial wetland cover, averaging 4.49 initial wetland acres per offset received, reflecting the fact that much of these parcels' area is initially classified as wetlands.

verify environmental improvements, which accrue over time. Table 1 shows that the regulator typically releases 15% of a wetland mitigation bank’s offsets once every three years, or an average time of $1/0.055 \approx 18.2$ years to build the entire project.

In contrast to wetland banks, the median developer purchases only 1.1 offsets (or 4.1 offsets on average). Measuring 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. Some development on wetlands involves on-site mitigation, and when offsets are purchased, the number of offsets per acre will vary based on the wetland value of the converted land. Fortunately, we use our ledger data to restrict the developed wetland pixel data to watersheds with active offset markets, in order to construct average acre-to-offset ratios for each water management district. Table A2 shows that, for markets with nonzero offset sales in a five-year period, wetland acres typically convert to offsets at a ratio of between 4.8 to 12.6 acres per offset.¹⁶

2.4.4 Market structure

The final aspect of the market design which—in combination with the initial land use conditions and the offset certification scheme described above—determines equilibrium trading is the designation of trading regions. As discussed above, banks can only trade within their own service areas, which approximate hydrological regions, and which we use to define markets after some adjustments to correct for partial and overlapping watersheds as discussed in Appendix B.

Figure 1 shows our market boundaries for Florida. On average, a market covers 1.15 million acres, or 33.5 watersheds, ranging from 11–70 watersheds. Wetland banks enter in 11.7% of market-years. The median market in the median year (2006–7) contained 1 bank or 2.2 on average, rising to 2 incumbents and an average of 3.7 firms by 2018. The average bank owns 26.1% of its market’s total production potential and 37.4% of its market’s total unsold offsets. The latter reflects the fact that banks do not sell their offsets immediately, but rather typically hold positive reserves; the median bank holds 52% of its offsets in reserve, with an interquartile range of [18%, 82%].

We interpret the market concentration in offset supply as reflecting a combination of the economies of scale and production delays discussed above, as well as strategic

¹⁶For a subset of transactions in the South Florida Water Management District, we observe parcels that purchased wetland offsets, which we link to our ledger of trades and intersect with our within-pixel calculations. For these parcels, developers purchase offsets at a ratio of 1 offset to 5.36 developed wetland acres on average, corroborating our water management district calculations.

factors that interact with the regional restrictions on trade. Conditional on various determinants of demand, annual entry occurs less often in markets with more incumbents (Table 4), consistent with incumbency advantages that deter future entry.

2.4.5 Trade outcomes

We close our discussion of the data with three descriptive findings about trade outcomes. First, offset markets have played an increasingly central role in the regulation and management of Florida’s wetlands since 1995. Valued at average annual prices, cumulative offset sales totaled \$1.2 billion from 1995–2018 (in 2020 USD). Offset sales also increase considerably over time, growing annually by an average of 9.8%. This growth reflects both a secular trend in new development in Florida and an underlying transition to the market-based approach to wetland conservation after the introduction of wetland mitigation bank rules in 1994.

Second, offset prices considerably exceed the observed components of banks’ average fixed costs, indicating the possibility of large markups.¹⁷ Observed fixed costs involve permitting costs, restoration and property maintenance costs (put into escrow at the time of entry), and the opportunity cost of land use. Total observed costs for the average bank, \$5.3 million, or about \$24,000 per offset, reflect primarily restoration and land costs. Average restoration costs, which we observe for nearly two-thirds of banks, are \$7,000/offset, while land values obtained from the last reported transaction price (from tax assessments) average \$19,000/offset (\$9,000/offset). These differences appear to reflect local land prices as well as natural features determine the costs and feasibility of restoration across markets.¹⁸

Third, high-development watersheds see much greater increases in flood insurance claims than wetland bank watersheds. On average, high-development watersheds experienced \$800,000 in claims per year from 2016–2020, in contrast to average claims for wetland bank watersheds, which in fact decline in real terms from \$315,000 to \$160,000 (2020 USD) per year. This difference continues through the distribution: the median high-development watershed experiences more than an order of magnitude the annual damages of the median wetland bank watershed. Some, but not all, of these differ-

¹⁷Given that costs are occurred up-front and sales occur stochastically over decades, these markups do not necessarily reflect large expected discounted profits. See Table 4 for estimated markups.

¹⁸For 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).

ences reflect differential growth in total insured value, with the average value at risk in high-development watersheds rising nearly twice as fast as in wetland bank watersheds (growing from \$18 to \$500 million relative to \$10 to \$66 million).

Taken together, our data indicates trade flows are substantial and that wetland banks typically locate in different places than where most wetland development occurs, even within relatively small regional markets. However, several empirical questions remain. First, transaction volumes and prices are not sufficient statistics for the private gains from trade. Inframarginal buyers may have values significantly greater than market prices, while imperfect competition may allow banks to charge prices above their costs. Second, 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 which 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. Third, evaluating offsets’ effects on hydrological outcomes like flooding 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.

3 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 3.1). A regulator who aims to conserve various wetlands issues permits to ensure that offsets satisfy its conservation objective (Section 3.2). Small land developers obtain payoffs from developing existing wetlands and take offset prices as given (Section 3.3). Large producers undertake long-run restoration activities to obtain offsets from the regulator, which they can sell to land developers over time. These producers incur fixed entry costs, have zero marginal costs, and take time to build (Section 3.4). Incumbents simultaneously choose sales in each period in a Markov perfect equilibrium (Section 3.5). In this setting, entry follows a cutoff rule and dynamic trading strategies can be characterized as an optimal inventory problem (Section 3.6).

3.1 The conservation problem

A large hydrological region or “market,” m , 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 . As offsets cannot be traded across markets, we suppress subscripts m until we introduce our estimating equations in Section 4. 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 $\delta < 1$.

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; i.e., $w_{it} = w_{i,t+1} = 1$. Second, locations with wetlands can be developed into non-wetland property and sold; i.e., $w_{it} = 1$ and $w_{i,t+1} = 0$. Third, land without wetlands can be restored into wetlands; i.e., $w_{it} = 0$ and $w_{i,t+1} = 1$.

Wetlands have private and social values. The private costs and benefits of wetland conservation, development, and restoration accrue to landowners and wetland restoration firms, which we describe in detail below. The social value of wetlands arise through their contribution to environmental quality. Specifically, each location i endowed with some potential environmental value as a wetland, $v_i \in \mathcal{V}$, where \mathcal{V} is some set of attributes. These potential values differ across space, due to underlying environmental characteristics and biodiversity, as well as over time, with evolving climatic and demographic conditions, though we suppress time dependence to simplify notation.

Given the private payoffs of wetland conservation, development, and restoration, in each period t , landowners will decide future land use for the next period 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 mitigation 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 into new development almost never transition back to wetlands.

3.2 Offset market design

The offset mechanism we study takes the following form. The regulator’s legal mandate is to ensure No Net Loss in wetland value given their conservation priorities. These

priorities are defined in our context by the regulator, who maps individual wetland attributes, $v_i \in \mathcal{V}$, into a scalar number of offsets, $\tilde{v}_i \in \mathbb{R}$. The regulator’s measure of aggregate environmental quality is then

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

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, i.e.,

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

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

Regulated trade in environmental offsets to satisfy (2) involves two types of participants in each period t . 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 impact, \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 or not 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 mitigation bank and traded in any future period.

We note that the irreversibility of both development and restoration considerably simplifies the dynamic land use problem in our setting. This is because it allows 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,” or private land with $w_{i0} = w_{it} = 0$, who decide whether or not to invest to restore their land into newly conserved wetland. These potential buyers and sellers of offsets share some similarities but also differ in a few economically important ways, so we analyze each type’s decision in separate sections.

3.3 Demand for offsets

Landowners whose land initially includes wetland correspond to potential buyers of offsets. These potential developers must seek approval from the regulator before developing their land, which requires offsets based on their wetland's attributes.

We assume a competitive market for private land development with a continuum of landowners, indexed by i . Each landowner inhabits one of the local 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, developing a wetland yields an ex-ante value, $u(X_{ht}, \xi_{ht}; \theta)$, 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 . This ex-ante value corresponds 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, a landowner i who develops a wetland incurs a choice-specific idiosyncratic development cost, ϵ_{it1} , assumed to be an independently and identically distributed Type 1 Extreme Value (T1EV) shock across i and t . Alternatively, a landowner who owns an existing wetland and chooses to do nothing obtains ϵ_{it0} , also independently and identically distributed T1EV.

Without regulation, the ex-ante private value for a landowner who develops on wetlands in watershed 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, under the market design of Section 2.2, 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

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

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} is then

$$\begin{aligned} Q_t(P_t, W_t, X_t, \xi_t, \tilde{v}; \theta) &= \int_0^1 w_{i0} w_{it} \tilde{v}_i \mathbf{1}\{u(X_{ht}, \xi_{ht}; \theta) + \epsilon_{it1} - \epsilon_{it0} \geq \tilde{v}_h \theta_P P_t\} di \\ &= \sum_h \tilde{v}_h W_{ht} \frac{e^{\theta' X_{ht} - \tilde{v}_h \theta_P P_t + \xi_{ht}}}{1 + e^{\theta' X_{ht} - \tilde{v}_h \theta_P P_t + \xi_{ht}}}, \end{aligned} \quad (4)$$

where the second line follow 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 private wetland availability, given by $W_t = \{W_{ht}\}_h$.

The structure of the private landowner's decision above imposes some important limitations on our analysis. First, while private wetland owners have the same average development payoffs as others in their local watershed, which creates correlation across development decisions within each local watershed, landowners act independently from one another and take offset prices as given. These assumptions rule out coordinated development schemes across many parcels. They reflect the small size of these developers relative to one another and to the banks described in Section 2.4. Second, in our model, prospective developers arrive, then disappear if they choose not to develop. While the decision rule in (3) incorporates the net present discounted value of development conditional on development, it rules out more complicated forward-looking strategies by developers that incorporate the option value of future development. This restriction limits our analysis to the extent that individual developers delay development to obtain more favorable offset prices or choice-specific shocks.¹⁹

Despite these restrictions, our model of wetland development remains rich enough to capture some essential aspects of the economic setting. We emphasize two primary advantages. First, aggregate market demand is only locally linear, because a market contains many local watersheds (more than thirty on average), and each watershed h has its own average utility. The marginal buyer of offsets will be a convex combination of these local watersheds, each with their own predilections for development, so our estimates of the curvature of demand and consumer surplus under observed trade flows will primarily reflect variation across local watersheds in their revealed preference for developing wetlands, not the logit distributional assumption on idiosyncratic choice

¹⁹More generally, optimal development will contrast the payoff in (3) with the option value of deferring the development decision to the next period. Where strategic delay in response to evolving offset prices is a crucial feature, the option value of future development can be modeled through assumptions about developers' expectations about future offset prices and the evolution of switching costs and tastes over time, as in, for example, Scott (2013).

shocks of each individual landowner.²⁰

Second, the model exhibits dynamics both within local watersheds and at the market level. Development on wetlands affect future development opportunities in each local watershed h , because the extent of available wetlands for development in period t , defined as $W_{ht} = \int_{i \in h} w_{i0} w_{it} di$, will reflect the full history of land development. These local stocks of potentially developable wetlands 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 specify their evolution with first-order Markov processes, so that the cumulative distribution function of the vector (X_{t+1}, ξ_{t+1}) is given by 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. However, for identification in Section 4, we will need to restrict the process ξ_t ; in particular, we assume that unobserved demand shifters $\xi_t = \{\xi_{ht}\}_t$ do not persist over time, except through permanent hydrological region fixed effects.²¹

3.4 Supply of offsets

We now turn to the choice problem for wetlands restoration, which—in contrast to the dispersed development of wetlands across space—involves a few large restoration sites in each market. We therefore model offset supply as an imperfectly competitive, oligopolistic environment with a finite set of non-infinitesimal potential producers, indexed by their location, $f \in \{1, 2, \dots, F\}$. Formally, 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 terms of natural suitability for restoration as well as intrinsic production potential, \tilde{v}_f . In this and the next subsection, we describe the

²⁰Note also that our data allows us to observe the initial condition (privately-held wetland), so we can specify this choice as a binary decision, which means that some of the usual caveats in multinomial logit models (e.g., requiring independence of irrelevant alternatives) do not apply here.

²¹Alternative assumptions about unobserved demand, such as permanent or persistent unobserved heterogeneity, are not difficult to accommodate if they take some known form, such as fixed effects or an AR(1) process. Then one can estimate our equation (14) below after applying the appropriate panel transformation. Alternatively, if one can construct a control function for ξ from some observed components of X (an example here would be to use local development on non-wetlands as a proxy for the unobserved local development payoff), then one can allow a nonparametric distribution for ξ .

producer's problem and timing; in Section 3.6, we close the model by specifying the equilibrium price path for offset prices.

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

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

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 non-wetland 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

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

offsets to each production site f . Offsets are released over time until restoration is complete, i.e., 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 environmental activities 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 from selling offsets to developers. At the start of each period t , each incumbent f has a stock of available offsets $B_{ft} \geq 0$, which have been certified but not yet sold. Each incumbent f can sell up to this constraint, $q_{ft} \leq B_{ft}$. Restoration costs are paid upfront, so the marginal cost of producing offsets is zero. We also assume that transaction costs are negligible.

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 then each firm f obtains per-period profits,

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

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

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

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

3.5 Information and timing

We denote the market state vector at time t by

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

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

The timing in each period t is as follows. All potential and current offset producers observe the market state, s_t . One prospective entrant $f \in \mathcal{F}_t^c$ privately draws their fixed cost, $\kappa_{ft} \sim G_t(\cdot | \tilde{v}_f, \mathcal{F}_t^c)$, and decides whether or not 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 decisions are realized, entry costs are incurred, wetlands are developed in proportion to permits sold, and the state updates to s_{t+1} .

3.6 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

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

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

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

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

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

of B_{ft} , T_{ft} , \tilde{v}_f , and s_t . All potential entrants use a common entry strategy that takes the form of a conditional cut-off rule: the pure strategy prescribes entry if and only if

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

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

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

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

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, i.e., $\sum_{f \in \mathcal{F}_t} q_{ft} = Q_t$ for all t , as well as $\lim_{t \rightarrow \infty} \delta^t P'_t B_{ft} = 0$ for all f .^{22,23}

²²The last condition, $\lim_{t \rightarrow \infty} \delta^t P'_t B_{ft} = 0$ for all f , is a standard transversality condition for infinite-horizon economies that rules out arbitrage in arbitrarily distant future periods by ensuring wetland bank values remain bounded (Acemoglu, 2008).

²³A Markov perfect equilibrium in symmetric pure strategies should exist for this game by Doraszelski and Satterthwaite (2010, Proposition 2). The reason is that, after conditioning on the number 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). The only added complication is the dynamic trading decision in (11), which is isomorphic to a continuous investment choice with evolving support. Doraszelski and Satterthwaite (2010) allow for continuous investment choices like $\mathcal{Q}(s_t, B_{ft}, T_{ft}, \tilde{v}_f)$ provided that they are unique functions of the state, as we assume here.

4 Empirical strategy and estimation

The empirical strategy to identify and estimate the model of Section 3 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 4.1). 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 *et al.* (2007) and Pakes *et al.* (2007) for both incumbents and potential entrants, to recover the distribution of fixed costs across market states consistent with optimal entry (Section 4.2). Third, we estimate the local flood externalities of different wetlands using historical changes in wetland extent and realized flood insurance claims, adapting the identification argument in Taylor and Druckenmiller (2022) to our setting with some slight modifications to minimize concerns of misspecification (Section 4.3).

4.1 Demand for offsets

We first describe how we obtain local demand for development on wetlands given offset prices. These estimates are necessary to calculate buyers’ surplus from observed offset purchases, to predict the pattern of local development on wetlands under alternative market designs, and to obtain the value of entry in this market.

To tractably account for some spatial correlation across infinitesimal locations, we partition the model’s continuum of locations i into the finite number of local watersheds, indexed by h as in Section 2. For each watershed, we recover local demand for environmental offsets, and then aggregate to total market-level demand as in (4), using water management district acre-to-offset ratios, \tilde{v}_h , to convert developed wetland acres into offsets. Our data allows us to construct pixel transitions over five-year intervals, so we estimate demand at the watershed by five-year period level, with periods t given by the intervals 1996–2001, 2001–2006, 2006–2011, 2011–2016. In particular, we can calculate the share of development occurring on private wetlands, $\omega_{ht} = Q_{ht}/W_{ht}$, as the ratio of the area Q_{ht} of private wetlands in watershed h converted to developed land over period t to the area of potentially developable wetlands $W_{ht} = \int_{i \in h} w_{it} w_{i0}$ at the start of period t . Taking this observed share ω_{ht} as the conditional probability that

$i \in h$ develops a private wetland, we obtain the logit equation

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

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 management district fixed effects, flood zone designations, new development on non-wetlands, 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 the vector of 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 projects.

First, we calculate the average production capacity of historical entrants whose service areas contain h . Intuitively, greater sunk capacity due to historical entry should shift market prices downwards, all else equal, acting as a downward cost shifter. Importantly, these capacities are fixed at the time of entry and cannot be subsequently adjusted. Clearly realized capacities are endogenous; the key is that when we control for the market state information known by those entrants, the realized draws become excluded shifters of future costs (Berry and Compiani, 2023). Because banks produce offsets slowly (over an average of eighteen years), our sunk capacity instrument can remain relevant over long horizons. The primary concern is that entrants rely on private information about future unobserved components of local demand. Our conversations with mitigation bankers, however, indicate they use market 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. In addition to using the average price from banks in the same water management district but not the local watershed's own market, we construct our average historical entrant capacity instrument from banks whose service area does not contain the watershed, but which are in the same water management district.

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 depend on the availability of private land for sale and the connectivity of that land to existing conservation land. Specifically, for each period and watershed, we

construct the total public wetland over 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 3 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.²⁴ These instruments vary in strength, with own historical capacity as the strongest instrument, with a first-stage F statistic exceeding 115, even conditional on our diverse controls, though the Hausman and public conservation land instruments also meaningfully shift prices (49.8 and 8.3). In addition, these instruments shift prices in the way theory predicts: all else equal, markets with larger historical entrants, more historical entrants in neighboring markets, or greater public conservation land, each have lower prices.

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.

4.2 Restoration costs

The main identification challenge to recovering unobserved production costs is that banks may enter more often in some markets because their fixed costs in those markets are especially low, because entry in those markets is particularly profitable, or some combination of the two. Our estimates of local offset demand, combined with structure on the entry and trading games, allow us to identify fixed costs using the equilibrium conditions derived in Section 3.4. In the Markov perfect equilibrium, trading strategies and entry decisions are given by the functions \mathcal{Q} in (11) and ϕ_t in (13). We take the

²⁴Without the instrument, column (1) of Table 3 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.

two-step approach of Bajari *et al.* (2007). First, we estimate flexible entry and trading strategies as well as production functions for wetland mitigation banks. Second, we calculate implied flow payoffs and value functions for incumbents, which we use to identify the distribution of fixed costs: conditional on those payoffs, the model implies that remaining variation in observed entry decisions will reflect 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 local characteristics of remaining production sites, \mathcal{F}_t^c , as well as the market conditions s_{mt} of the broader market m . Our data includes the location and date of every wetland mitigation bank as well as ownership and land characteristics everywhere within the hydrological boundary of each market. First, we use this data to construct a measure of the extent of remaining land available for wetland mitigation restoration, which we proxy with a vector of market-level characteristics X_{mt} , which include the areas of public and private wetland and developed land. Second, we use the location and timing of entry to estimate the entry model given by (13). Specifically, we estimate annual market-level entry probabilities at the market-year level for all markets m and years $t \in \{1995, 1996, \dots, 2020\}$ according to the following probit specification,

$$\mathbb{P}(\{\text{enter}_{mt} = 1\} | X_{mt}, s_{mt}) = \Phi(\beta_1'(X_{mt}, s_{mt})), \quad (15)$$

where Φ is the Gaussian CDF, X_{mt} proxies for the extent of remaining production sites as described above, and the market state s_{mt} is the same as used in (14).

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 non-wetlands, and fewer incumbents. For example, Table 4 shows that the average estimated annual entry probability for duopoly markets is 0.18 for the first firm but only 0.11 for the second firm.

Production function. We observe the numbers and dates of offsets issued to each bank directly from various regulatory records, which correspond to b_{ft} for banks f and years t in our framework. We also observe each bank's total offset allowance for the lifespan 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 analogue of (6) in two pieces. First, we are interested in the total offset allowance \tilde{v}_f . Second, we are interested in the timing of offset releases

over time, 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 \hat{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 2.4.

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}^{sold} , the number of offsets issued to, and sold by, each bank. This lets us estimate trading strategies at the incumbent-year level from 1995–2018. An incumbent bank’s sales q_{ft}^{sold} in year t depend on its reserves and future production, as well as characteristics of its competitors and demand, via

$$q_{ft}^{\text{sold}} = \chi(s_{-ft}, B_{ft}, T_{ft}, \tilde{v}_f, s_t), \quad (16)$$

where χ is a flexible polynomial, rivals’ characteristics s_{-ft} include $N_{\text{competitors of } f}$ and $(B_{f't}, T_{f't})_{f' \text{ is a competitor of } f}$, and the market state s_t includes demand characteristics $\{X_{ht}\}_{h \in m}$ from (14) and private wetland stocks $\{W_{ht}\}_{h \in m}$ of watersheds in market m .

As in many applications of Bajari *et al.* (2007), the policy function χ consistent with the model is a nonparametric function of a high-dimensional state space, so its estimation in a finite sample may lead to error. In our simulations, the rules that allow banks only to trade certified offsets significantly limit these concerns by bounding $\chi \in [0, B_{ft}]$. We further discipline χ by imposing individual rationality (IR) constraints derived from static Cournot first-order conditions using the aggregate (market) demand elasticity and the vector of equilibrium market shares; this proves quite useful in practice.

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 non-wetlands depends significantly on the previous stock of developed land, and population and income are highly persistent, with an autocorrelation coefficient of 0.97. 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

$$V(B_{f0}, T_{f0}, s_0, \tilde{v}_f) = \sum_{t=0}^T \delta^t \int_{S^t} \Pi(\mathcal{Q}(s_t, B_{ft}, T_{ft}, \tilde{v}_f), s_t) dH^t(s_t | s_0) \quad (17)$$

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 . The transition from s_t to s_{t+1} updates the ledger to account for trading and production, then draws new entrants and demand shocks. The ledger moves from B_{ft} to $B_{f,t+1}$ according to (8), with q_{ft} given by (11) and new production b_{ft} given by (6) for each $f \in \mathcal{F}_t$. Then, entry in each location $f \notin \mathcal{F}_t$ occurs with probability $\phi_t(s_t, \mathcal{F}_t^c)$. We add the new entrants, if any, to the list of incumbents \mathcal{F}_{t+1} . Finally, the demand shifter evolves from (X_t, ξ_t) to (X_{t+1}, ξ_{t+1}) , with X_{t+1} drawn from $H_X(\cdot | X_t)$, the AR(1) state transition estimates discussed above and ξ_{t+1} drawn from H_ξ , the empirical distribution of residuals.

From realized trading volumes $Q_t = \sum_f q_{ft}$ for all $t \geq 0$, we obtain prices $(P_t)_{t \geq 0}$ from (4) for the sequence of realized entry. This then allows us to calculate $\Pi(\cdot)$ with (7) for each t along this sample path. With H obtained above, the distribution of these payoffs allows us to obtain the value function in (17). 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 $G_t(\cdot | \tilde{v}_f, \mathcal{F}_t^c)$. Appendix C describes the algorithm in detail.

Estimates. Table 4 reports results for our entry cost estimator, as well as structural parameters. Conditional on entry, average fixed costs over entrants are estimated to be \$6.6 million, or \$21,500 per offset certified (median \$13,500), with considerable dispersion across banks, with interquartile range of \$1,440 to \$32,000/offset.

Importantly, the estimated costs resemble observed costs discussed in Section 2.4 but not used in estimation. Table 4 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, other than unobserved permitting costs, the two major costs of wetland banking should be restoration and the opportunity cost of land.

The structural parameters reported in Table 4 also provide some additional insight into entry costs. First, the 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 mitigation banking. Second, the estimated markups and realized rates of return on capital appear plausible, averaging 7.8% and

with an interquartile range of 3.7–10.4%, which are comparable to the average real rate of return of 5.86% on U.S. housing from 1980–2015 (Jorda *et al.*, 2019, Table 7).

4.3 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.

Identifying flood protection benefits. 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 our discussion of the regulations and economic incentives for land use emphasized, wetlands are not randomly assigned. The primary threats to identification are changes in land use that (a) heighten exposure to flood risk and (b) correlate with—but are not caused by—changes in wetland extent. Therefore, we control for historical insured flood claims and the prior extent of developed land in each local area, as well as the local watershed’s permanent underlying flood risk as measured by flood hazard maps. In addition, we observe the source of new development using state transitions for each pixel, which allow us to calculate the share of development on wetlands separately from the share of new development on other, non-wetland vacant land. This allows us to control directly for new development on vacant land. We view this control as a proxy for some of the unobserved shocks to land development payoffs that correlate with both wetland destruction and underlying exposure to flood risk. Finally, as an outcome, we use only flood insurance claims for structures that were built prior to 1995, to ensure that our measure primarily reflects the spillovers from wetland protection, not the new properties built on wetlands that are (mechanically) exposed to flood.

We assume a local Cobb-Douglas flood damage function $D_h(\cdot)$, where wetlands provide flood protection through a constant elasticity and proportional to the underlying risk of the local watershed. The exponential functional form is particularly important given that observed flood damages range over eight orders of magnitude across watersheds (Table 1). In addition, converting wetlands to development and restoring wetlands as banks need not have symmetric effects on flood damages. Development on

wetlands, Q_{ht} , changes damages with an elasticity ζ_d ,

$$\frac{\partial}{\partial Q_{ht}} D_h(W_{ht}, X_{ht}, Q_{ht}, B_{ht}) = \zeta_d Q_{ht}^{\zeta_d-1} B_{ht}^{\zeta_b} D(X_{ht}, W_{ht}) e^{\varepsilon_h} \quad (18)$$

whereas the restoration of wetland mitigation banks, B_{ht} , involves marginal changes of

$$\frac{\partial}{\partial B_{ht}} D_h(W_{ht}, X_{ht}, Q_{ht}, B_{ht}) = \zeta_b B_{ht}^{\zeta_b-1} Q_{ht}^{\zeta_d} D(X_{ht}, W_{ht}) e^{\varepsilon_h}. \quad (19)$$

Our baseline specification to estimate D , ζ_d , and ζ_b uses ex-post local outcomes, average annual flood damages in the post-period (2016–20) across local watersheds h , to study wetland changes under the offset market mechanism observed from 1996–2016. We use flood damages in the pre-period (1991–1995) to control for unobserved confounders, with our preferred specification being an AR(1), though we find similar elasticities with a long first difference, which assumes a unitary persistence but identifies the elasticities under a broader range of unobservables. The estimating equation is

$$\begin{aligned} \text{asinh}(\text{claims}_{h,\text{post}}) &= \zeta_d \ln Q_{h,1996-2016} + \zeta_b \ln B_{h,1996-2016} \\ &\quad + \rho \cdot \text{asinh}(\text{claims}_{h,\text{pre}}) + \gamma' X_h + \varepsilon_h, \end{aligned} \quad (20)$$

where $\zeta = (\zeta_d, \zeta_b)$ are the coefficients of interest for development on former wetlands $Q_{h,1996-2016}$ and newly-created wetlands, $B_{h,1996-2016}$, and X_h includes new development on non-wetlands, percent area in baseline flood risk categories (A and V zones), wetland and developed land areas at baseline, and water management district fixed effects.

Note that (20) imposes four simplifying assumptions on wetland flood protection benefits. First, we follow prior literature to assume that lost wetlands affect floods through their extent or acreage (Brody *et al.*, 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 elasticity in (20) implies that level differences in local protection values arise through the intercept, based on differences in historical exposure, via $\rho \cdot \text{asinh}(\text{claims}_{h,1991-1995})$, and other local characteristics, such as the extent of developed land and baseline flood hazard risk, via $\gamma' X_h$. Third, we estimate (20) at the local watershed level, which captures the within-watershed externalities of local development for 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, which indicates that the local watershed is an appro-

priate spatial unit of analysis for our study.²⁵ Fourth, floods involve economic damage beyond insurance claims. For example, our measure will not account for flood damage to uninsured properties, damage to insured properties that exceed policy limits, or the cost of defensive investments undertaken to lower flood risk.²⁶

Estimates. Table 5 presents the results of estimates of (20) across different controls and subsamples. Column (1) shows a strong positive relationship between the amount of development on wetlands and the total claims for flood insurance damages, which is consistent with the prior literature’s findings that wetlands reduce flood damage. Interestingly, wetland bank area positively correlates with flood damage as well, despite the array of controls, which include contemporaneous development on land other than wetlands, as well as baseline developed land and flood zones.

Moving from column (1) to the subsequent columns, however, highlights the severity of the omitted variable bias that arises without controlling for prior insured damages. This omitted variable bias seems likely to bias upwards the wetland development coefficient, ζ_d , if places with considerable development on wetlands between 1996–2016 are places with growing exposure to flood risk. It also seems likely to bias upwards (and perhaps reverse the sign) of the wetland bank area coefficient, given that wetland banks enter in places where they can serve developers interested in buying offsets.

Columns (2)–(6), which add the control for baseline flood claims, significantly reduce the implied elasticity wetland-development elasticity. The persistence coefficient is precisely estimated at 0.42 and leads to a much better fit, increasing the adjusted R^2 to 0.53 from 0.43. The elasticity of flood damages with respect to wetland development falls to 0.267, about half of its value without this control (and less than one-fifth of the naïve estimate without controls). This estimate is still significant at the 1% level and, as we show below, substantial in magnitude. Furthermore, (2) shows that the correlation between land committed to wetland mitigation banks and local flood damages in the OLS no longer appears statistically significant from zero.

Column (3) repeats (2) adding demographic controls and (4) adds hydrological region fixed effects; the damage and persistence coefficients are not distinguishable from

²⁵See Table A5 for more details. A watershed’s location in the hydrological network mildly predicts its flood damages. However, wetland development in neighboring watersheds do not predict local damages. Interestingly, upstream wetland bank activity is correlated with local damages; this appears to reflect the correlation between banking and local wetland development, as the total damages with this richer damage function are indistinguishable from our main specification.

²⁶In 2015, 38.7% of Florida households (52.5% 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% of total claims (2.9% of building claims and 15% of content claims).

the original estimates. Columns (5) and (6) show a similar elasticity obtains when we take first differences and when drop the watersheds that do not have flood insurance in 1995. For parsimony, we use (2) as our preferred specification when we evaluate the effects of wetland reallocation on insured flood damages below.²⁷ Furthermore, given the insignificant estimates of wetland bank area in specifications (2)–(6), reported in the second row of Table 5, we assume below that wetland banks themselves do not directly affect flood outcomes, though this is easily relaxed.

Our flood protection estimates compare favorably with some recent work on flood protection and wetlands, which we summarize in Table A6. They imply annual flood damage spillovers from development on Florida wetlands of about \$1,400 per hectare on average, which resemble earlier studies finding average annual wetland flood protection values in the Gulf Coast that translate to \$511 (2020 USD) per hectare in Florida (Brody *et al.*, 2015). For high-risk storm flood zone watersheds, we estimate annual flood damages of \$25,200 per hectare, not dissimilar from recent estimates of \$18,000 per hectare 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.²⁸ The order-of-magnitude discrepancy between our results and theirs likely arise from specification differences. TD 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 ($R^2 \approx 0.055$). We specify a nonlinear model which 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 ($R^2 \approx 0.531$). We also take a different approach to measuring spillovers than TD (damage to structures built before 1995, not all damages in neighboring zip codes) because it appears to better explain our data (footnote 25).

²⁷The flood protection function estimates in Table 5 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 (Table A3).

²⁸Taylor and Druckenmiller (2022) estimate annual causal effects of \$12,081 per hectare of wetlands converted to development and \$8,290 per hectare of wetland lost in highly developed areas from 2001–16. In Table A6, we calculate that 49,700 hectares of wetland were converted into development and 56,700 hectares of wetland disappeared in highly developed areas from 2001–2016, while total annual flood claims increased by about \$270 million in Florida and \$144 million in its highly-developed areas, or \$5,400 and \$2,530 per hectare, respectively. That is, the Taylor and Druckenmiller (2022) estimates imply that observed wetland changes over this period should have caused 223% and 327% of the observed increases in flood claims, respectively. Our model, in contrast, attributes 26% and 24% of observed post-2015 flood claims to wetland development in offset markets.

5 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. Specifically, we contrast the observed reallocation under the Florida offset market mechanism from 1995–2018 with two sets of counterfactuals. First, we evaluate the market relative to historical conservation rules, in order to assess the private gains from trade (Section 5.1) and flood externalities (Section 5.2) from the transition to the market-based mechanism. Second, we analyze ways to improve the design of the offset market (Section 5.3), given our new estimates of private gains from trade and flood externalities.

5.1 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,

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

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

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

Figure 3B plots consumer surplus over all trades, ordered by descending average surplus and weighted by the number of offsets.

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 aggregate producer surplus, PS_m , by

building an aggregate marginal producer surplus curve that we integrate over observed trades, as described in Appendix C. Figure 4B plots realized producer surplus and costs, ordered by descending producer values.

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

$$\text{GFT}_m = \text{CS}_m + \text{PS}_m$$

Table 7 reports the results for all of Florida, $\sum_m \text{GFT}_m$, our first key empirical finding. The first column shows estimates of developer values, bank costs, and the private gains from trade under the market. The first row shows wetland development under the market under our simulation, which closely resembles actual wetland development from 1996–2016. The second row shows aggregate developer valuations for these trades, i.e., $\sum_t U_{mt}$, which equal about \$2.8 billion (2020 USD). Total fixed costs, reported in the next row, are about \$600 million, implying private gains from trade of about \$2.2 billion. Given total sales (\approx \$1.1b), consumer surplus from the demand estimates from Section 4.1 equals \$1.7 billion, while producer surplus is about \$500 million. These estimates indicate that the private gains from offset trade accrue to both developers of wetlands and wetland banks.

5.2 Flood externalities

We now construct marginal environmental damages using our location-specific panel estimates for insurance payouts. Given that development of wetlands is irreversible, the social cost of removing flood protection benefits corresponds an infinite sequence of discounted damages; we therefore scale our annual average effect by $\sum_{t=0}^{\infty} 0.95^t$ using a real discount rate of 5% in accordance with federal government regulatory guidelines, though we also report totals for 3% and 7%. We can then obtain marginal damages given by (18) from our estimates of $\hat{\zeta}_d$, $\hat{\zeta}_b$, $\hat{\gamma}$, and $\hat{\rho}$ from Table 5 as well as the data on historical claims and other observables at baseline.²⁹

Table 6 reports the distribution of the local flood protection estimates of wetlands across watersheds with wetland development under the offsets market. The externality from developing a wetland in the median watershed is \$1,110/offset, which is a rounding

²⁹While we calculate the direct effect of development on post-period flood damages, we do not solve for long-run damages with the full lag structure implied by (20). That is, we interpret the lag in (20) as a control for baseline risk, rather than a causal model where past damages increase future damages.

error from the viewpoint of a land developer, given a typical price of \$80,000/offset. Hence for many watersheds, wetlands’ local flood protection benefits do not justify altering trading rules. However, the highest-percentile externalities (e.g., 90%, 95%-ile, of \$49,500/offset and \$115,550/offset) are comparable to observed offset prices. This dispersion is also clear from Figure 6, which plots estimated flood damages for each development occurring on wetlands from 1996–2016. The jagged blue peaks show high risks in some places amidst many wetlands that deliver little or no flood protection value. Figure 7, Panel A overlays these estimates with each project’s private value. Wherever the blue spikes cross the red line, development occurred despite estimated flood benefits of conservation that exceed developer values.

Integrating damages over all development on wetlands in each offset market, we can approximate the total flood damage arising from wetland offset trading over time. This is our second major empirical finding. We find wetlands whose disappearance we attribute to offset trade from 1996–2016 would have delivered \$1.0–1.6 billion of flood protection, depending on whether outliers (watersheds above the 99.9%, 99% and 97.5%-ile, 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, Table A3 reports marginal and total damages for some alternative estimates of (20). Both the distribution of flood protection values across wetlands and the total excess flood damage appear similar to our baseline specification, though the tails above 99% appear to be sensitive to the definition of historical flood exposure.

5.3 Optimal policy

Using the results from Table 6, we can approximate the optimal Pigouvian taxes with the marginal damage of development at h . These taxes are differentiated across location based on the observables in (20). We further simplify the counterfactual analysis in two helpful but restrictive ways. First, we use our aggregate cost function to avoid the need to solve for the new dynamic equilibrium to determine costs. Second, we rule out dynamic interdependencies in demand and new equilibrium prices by evaluating each period’s demand with the tax applied to observed prices, rather than iterating the state variable forward and simulating equilibrium bank trading strategies. This is straightforward to relax by forward-simulating demand, so that counterfactual demand

in earlier periods alters future demand through wetland and developed land states.

These simplifications allow us to evaluate the vector of Pigouvian taxes directly with our existing demand, cost, and externality estimates. The second column of Table 7, which contains the results, is the third major empirical finding of our paper. Introducing a simple modification of trading rules to account 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 optimal Pigouvian design creates an average of two dollars of flood protection benefits for each dollar of gains from trade forgone ($\frac{1605-258}{2261-1593} = \frac{1347}{668} = 2.03$). Crucially, the design maintains No Net Loss; 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 development in Florida. This policy is of economic interest for at least two reasons. First, comparing the local Pigouvian design with a uniform tax illustrates the extent to which 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 tax 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 corrective tax that accomplishes this objective turns out to be \$30,000/offset, or about one-third of the mean price through the sample. As the third column of Table 7 shows, such a policy significantly lowers flood damages relative to the market, but at a much higher private cost. The uniform tax preserves less than half of the market’s private gains from trade and requires a significant decline in development, lowering development on wetlands to 107,309 acres relative to the 140,650 acres under the market and 82,000 acres under the Pigouvian design. As Figure 7, Panel C shows, despite the significant reduction in development on wetlands in this counterfactual, estimated damages significantly exceed the private surplus for much of the development that occurs, underscoring the need for policy that can target local watersheds based on underlying flood risk.

6 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. We see our approach as applicable to a broad range of environmental markets where the regulator accesses data on offsets 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 useful for analyzing markets for conservation offsets in particular, as well as other markets for environmental offsets where the production technology for offsets differs from the cost of direct 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 some important policy implications for the \$1 trillion Florida economy, where real estate accounts for nearly 20% of GDP and wetlands comprise 29% of land, and the qualitative results may generalize to wetland offset markets beyond Florida. 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 welfare-improving policy holding fixed the regulator’s existing conservation objectives. A Pigouvian correction based on observable local characteristics lowers excess flood damages by more than 80% 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 one-eighth of the flood reduction benefits of the local Pigouvian design. We view the robustness of these empirical findings as a key area for future research.

References

- ACEMOGLU, D. (2008). *Introduction to Modern Economic Growth*. Princeton University Press.
- ARMY-EPA (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**, 9210.
- ARONOFF, D. and RAFEY, W. (2020). *Essays in Environmental Markets*, Ph.D. Dissertation, Department of Economics, MIT, chap. 2. Conservation priorities and environmental offsets.
- ASKER, J., COLLARD-WEXLER, A. and DE LOECKER, J. (2019). (Mis)allocation, market power, and global oil extraction. *American Economic Review*, **109** (4), 1568–1615.
- ASSUNÇÃO, J., MCMILLAN, R., MURPHY, J. and SOUZA-RODRIGUES, E. (2019). Optimal environmental targeting in the Amazon Rainforest. *NBER Working Paper*, **25636**.
- BAJARI, P., BENKARD, C. L. and LEVIN, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, **75** (5), 1331–1370.
- BEA (2020). *Annual Gross Domestic Product (GDP) by state: Real estate and rental and leasing*. Tech. Rep. SQGDP2N, Bureau of Economic Analysis.
- BERGSTROM, T., BLUME, L. and VARIAN, H. (1986). On the private provision of public goods. *Journal of Public Economics*, **29**, 25–49.
- BERRY, S. and COMPIANI, G. (2023). An instrumental variable approach to dynamic models. *Review of Economic Studies*, **90** (4), 1724–1758.
- BORENSTEIN, S. *et al.* (2002). Measuring market inefficiencies in California’s restructured wholesale electricity market. *American Economic Review*, **92** (5), 1376–1405.
- BOYD, J. and WAINGER, L. (2002). Landscape indicators of ecosystem service benefits. *American Journal of Agricultural Economics*, **84** (5), 1371–1378.
- BRODY, S., HIGHFIELD, W. and BLESSING, R. (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), 1556–1567.
- CARLSON, C. *et al.* (2000). Sulfur dioxide control by electric utilities: What are the gains from trade? *Journal of Political Economy*, **108** (6), 1292–1326.
- DIAMOND, P. (1973). Consumption externalities and imperfect corrective pricing. *Bell Journal of Economics*, **4**, 526–538.
- DORASZELSKI, U. and SATTERTHWAIT, M. (2010). Computable Markov-perfect industry dynamics. *The RAND Journal of Economics*, **41** (2), 215–243.
- ERICSON, R. and PAKES, A. (1995). Markov-perfect industry dynamics: A framework for empirical work. *Review of Economic Studies*, **62** (1), 53–82.
- ERWIN, K. (2009). Wetlands and global climate change: The role of wetland restoration in a changing world. *Wetlands Ecology and Management*, **17**, 71–84.
- FEMA (2018). *An affordability framework for the National Flood Insurance Program*. Tech. rep., Federal Emergency Management Agency.
- FERNANDEZ, L. and KARP, L. (1998). Restoring wetlands through wetlands mitigation banks. *Environmental and Resource Economics*, **12**, 323–344.
- FLORIDA STATE LEGISLATURE (2019). Title XXVIII. Natural resources; conservation, reclamation, and use. chapter 373. Water resources.

- FOWLIE, M. and MULLER, N. (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.
- , REGUANT, M. and RYAN, S. (2016). Market-based emissions regulation and industry dynamics. *Journal of Political Economy*, **124** (1), 249–302.
- HAUSMAN, J. (1996). *The Economics of New Goods*, University of Chicago Press, chap. Valuation of new goods under perfect and imperfect competition, pp. 209–248.
- HERNANDEZ-CORTES, D. and MENG, K. (2020). Do environmental markets cause environmental justice? Evidence from California’s carbon market. *NBER Working Paper*, **27205**.
- HODGSON, C. (2018). Information externalities, free riding, and optimal exploration in the UK oil industry. *Job Market Paper, Stanford University*.
- HOTZ, V. J., MILLER, R., SANDERS, S. and SMITH, J. (1994). A simulation estimator for dynamic models of discrete choice. *Review of Economic Studies*, **61** (2), 265–289.
- HSIAO, A. (2021). Coordination and commitment in international climate action: Evidence from palm oil. *Job Market Paper, MIT*.
- JACK, K., KOUSKY, C. and SIMS, K. (2010). Designing payments for ecosystem services. *Proceedings of the National Academy of Sciences*, **105**, 9465–9470.
- JOHNSTON, R. *et al.* (2002). Combining economic and ecological indicators to prioritize salt marsh restoration actions. *American Journal of Agricultural Economics*, **84** (5), 1362–1370.
- JORDA *et al.* (2019). The rate of return on everything, 1870–2015. *Quarterly Journal of Economics*, **134** (3), 1225–1298.
- KOTCHEN, M. (2009). Voluntary provision of public goods for bads: A theory of environmental offsets. *The Economic Journal*, **119** (537), 883–899.
- KOUSKY, C. and WALLS, M. (2014). Floodplain conservation as a flood mitigation strategy: Examining costs and benefits. *Ecological Economics*, **104**, 119–128.
- LUPI, F. *et al.* (2002). The economic equivalency of drained and restored wetlands in Michigan. *American Journal of Agricultural Economics*, **84** (5), 1355–1361.
- MASKIN, E. and TIROLE, J. (2001). Markov perfect equilibrium: I. Observable actions. *Journal of Economic Theory*, **100** (2), 191–219.
- MCCOMBS, J. *et al.* (2016). Accuracy assessment of NOAA Coastal Change Analysis Program 2006–2010 land cover and land cover change data. *Photogrammetric Engineering & Remote Sensing*, **82** (9), 711–718.
- MULLER, N. and MENDELSON, R. (2009). Efficient pollution regulation: Getting the prices right. *American Economic Review*, **99** (5), 1714–1739.
- PAKES, A., OSTROVSKY, M. and BERRY, S. (2007). Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). *The RAND Journal of Economics*, **38** (2), 373–399.
- POLASKY, S. (2002). The economics of wetland ecosystem restoration and mitigation: Discussion. *American Journal of Agricultural Economics*, **84** (5), 1379–1380.
- RAFEY, W. (2023). Droughts, deluges, and (river) diversions: Valuing market-based water reallocation. *American Economic Review*, **113** (2).
- RAMSAR CONVENTION (2018). *Global Wetland Outlook: State of the World’s Wetlands and their Services to People*. Tech. rep., Ramsar Convention on Wetlands, Gland, Switzerland.
- RAPANOS V. UNITED STATES (2006). 547 U.S. 715.

- RUHL, J. *et al.* (2009). Implementing the new ecosystems services mandate of the Section 404 compensatory mitigation program. *Stetson Law Review*, **38**, 251–272.
- SAIZ, A. (2010). The geographic determinants of housing supply. *Quarterly Journal of Economics*, **125** (3), 1253–1296.
- SALZMAN, J. and RUHL, J. (2006). *Moving to Markets in Environmental Regulation*, Oxford University Press, chap. 12. “No net loss”: Instrument choice in wetlands protection.
- SAMUELSON, P. (1976). Economics of forestry in an evolving society. *Economic Inquiry*, **14** (4), 466–492.
- SCOTT, P. (2013). Dynamic discrete choice estimation of agricultural land use. *Working paper*, Princeton University.
- SILVERSTEIN, J. (1994). Taking wetlands to the bank. *Boston College Environmental Affairs Law Review*, **22**, 129–161.
- SMALL, K. and ROSEN, H. (1981). Applied welfare economics with discrete choice models. *Econometrica*, **49** (1), 105–130.
- SOUZA-RODRIGUES, E. (2019). Deforestation in the Amazon: A unified framework for estimation and policy analysis. *Review of Economic Studies*, **86** (6), 2713–2744.
- STAVINS, R. and JAFFE, A. (1990). Unintended impacts of public investments on private decisions: The depletion of forested wetlands. *American Economic Review*, **80** (3), 337–352.
- SUN, F. and CARSON, R. (2020). Coastal wetlands reduce property damage during tropical cyclones. *Proceedings of the National Academy of Sciences*, **117** (11), 5719–5725.
- SUNDING, D. and ZILBERMAN, D. (2002). The economics of environmental regulation by licensing: An assessment of recent changes to the wetland permitting process. *Natural Resources Journal*, **42**, 59–90.
- TAYLOR, C. and DRUCKENMILLER, H. (2022). Wetlands, flooding, and the Clean Water Act. *American Economic Review*, **112** (4), 1334–1363.
- TURNER, M., HAUGHWOUT, A. and VAN DER KLAAUW, W. (2014). Land use regulation and welfare. *Econometrica*, **82** (4), 1341–1403.
- USGS (2013). *Federal standards and procedures for the National Watershed Boundary Dataset (WBD), Fourth edition*. Tech. rep., U.S. Geological Survey and U.S. Department of Agriculture, Natural Resource Conservation Service.
- WING, O. *et al.* (2018). Estimates of present and future flood risk in the conterminous United States. *Environmental Research Letters*, **13**, 034023.

TABLE 1. NEW DATA ON WETLAND OFFSETS IN FLORIDA

	N	avg	sd	q0	q25	q50	q75	q100
Initial wetland (pct/pixels)	136,302,645	36.5	48.1	0	0	0	100	100
Initial developed land (pct/pixels)	136,302,645	12.0	32.5	0	0	0	0	100
<u>Hydrology</u>								
Area ('000 pixels/watershed)	1,004	135.8	168.6	27.7	79.0	109.7	146.1	2,753.8
Area ('000 acres/watershed)	1,004	30.2	37.5	6.2	17.6	24.4	32.5	612.4
<u>Land Ownership</u>								
Private land ('000 acres/watershed)	1,004	25.3	23.9	0	14.7	20.9	29.9	388.9
Public land ('000 acres/watershed)	1,004	4.8	27.5	0	0	0.5	3.8	564.2
Percent public land (watershed)	1,004	12.5	21.8	0	0	2.1	14.7	100
<u>Initial Conditions</u>								
Initial wetlands ('000 acres/watershed)	1,004	11.0	30.6	0.001	4.1	7.2	11.6	542.3
Initial wetlands (pct/watershed)	1,004	34.0	20.1	0.002	20.0	30.8	44.0	99.6
Initial private wetlands ('000 acres/watershed)	1,004	7.7	13.8	0	3.2	5.7	9.3	387.3
Initial public wetlands ('000 acres/watershed)	1,004	3.3	25.0	0	0	0.3	1.9	528.1
Pct public wetland (pct/watershed)	1,004	17.2	24.7	0	0	4.8	25.0	100
Initial developed land ('000 acres)	1,004	3.6	5.8	0	0.4	1.2	4.5	49.7
Initial developed land (pct)	1,004	13.8	19.1	0	1.8	4.7	18.2	92.8
<u>Wetlands Development and Restoration</u>								
Wetlands developed, 1996-2016 (acres/watershed)	1,004	207.5	483.4	0	2.4	16.3	186.7	4,908.3
Private wetlands developed, 1996-2016 (acres/watershed)	1,002	206.2	481.3	0	2.0	15.9	186.1	4,812.2
Public wetlands developed, 1996-2016 (acres/watershed)	731	2.4	8.4	0	0	0	1.1	96.1
$\mathbb{P}(\text{develop} \text{wet}) \times 100^a$	1,000	3.7	7.3	0	0.04	0.3	3.9	57.0
$\ln \frac{\mathbb{P}(\text{develop} \text{wet})}{1 - \mathbb{P}(\text{develop} \text{wet})}$	912	-5.3	2.6	-12.0	-7.3	-5.4	-3.0	0.3
Area of wetland banks, 2016 (acres/watershed)	1,378	132.2	663.9	0	0	0	0	9,518.7
<u>Market Structure</u>								
Area ('000 acres/market)	30	1,153.3	800.4	264.0	516.8	863.4	1,520.4	3,993.5
Area (watersheds/market)	30	33.5	18.0	11	18.2	27.5	49	70
Initial private wetlands ('000 acres/market)	30	257.5	215.6	43.5	100.1	188.5	355.8	912.9
Entry ^b (market-year)	780	11.7	32.1	0	0	0	0	100
Number of banks (market-year)	530	2.6	2.3	1	1	2	3	16
Lagged capacity ^c (credits/market-year)	780	556.9	1,214.7	0	0	0	535	7,137
Annual trades (credits/market-year)	381	41.7	67.3	0.02	2.9	11.7	50.6	516.0

Descriptive statistics for Florida, 1995-2020.

^a $\mathbb{P}(\text{develop}|\text{wet})$ defined as the within-pixel probability that a wetland pixel in 1996 becomes developed in 2016.^bWetland bank entry indicator equals 1 if a wetland mitigation bank enters in market m in year t and 0 otherwise.^cHistorical capacity defined as lifetime offset production authorized to all banks in market m that entered $k = 5$ or more years prior to t .

TABLE 1 (cont'd): NEW DATA ON WETLAND OFFSETS IN FLORIDA

	N	avg	sd	q0	q25	q50	q75	q100
Demand Shifters								
Other land developed, 1996-2016 (acres/watershed)	1,004	518.9	985.9	0	43.8	142.4	522.9	9,400.3
Average annual home price ('000\$ per watershed)	999	143.2	65.2	47.2	98.4	135.6	172.8	868.0
Endogenous State Variables								
Wetlands stock ('000 acres/watershed/period)	5,020	10.8	30.6	0.001	4.0	6.9	11.4	542.5
Private wetlands ('000 acres/watershed/period)	5,020	7.5	13.7	0	3.1	5.6	9.1	387.5
Public wetlands ('000 acres/watershed/period)	5,020	3.3	25.0	0	0	0.3	1.9	528.1
Developed land stock ('000 acres/watershed/period)	5,020	4.1	6.4	0	0.5	1.4	5.4	56.5
Offset Credit Release and Sales								
Bank entry year	107	2,008.1	7.5	1,995	2,003	2,009	2,014.5	2,020
Bank size (acres/bank)	107	1,866.1	2,680.0	66	428.5	1,049	2,157.5	22,805
Bank size (credits/bank)	106	410.0	566.1	13	85.2	203	521.8	4,345
Acres per credit	106	5.9	4.5	1.1	3.1	5.1	6.9	29.4
1(credits released) per bank per year	1,209	0.3	0.4	0	0	0	1	1
Credits released per bank per year (pct total)	343	15.3	16.2	0.05	5	10.0	20.0	96.8
Annual sales (credits/bank-year)	981	15.5	31.4	0	0	1.8	15.4	236.4
Bank reserves ^d (pct/bank-year)	967	51.8	33.6	0	18.3	54.7	82.0	100
Acre wetland developed per credit sold	5,512	8.8	2.8	4.8	8.1	8.2	11.6	12.6
Credit price ('000\$/credit), all transactions	1,432	87.5	61.7	1.0	38.6	63.4	137.2	785.3
Credit price ('000\$/credit), average per market per year	151	98.8	50.5	5.5	62.0	93.9	127.2	306.2
Credit price ('000\$/credit), average per market per year, sd	61	51.2	36.5	1.0	24.3	49.3	73.0	203.0
Flood Risks								
Flood zone (pct/watershed)	1,004	41.7	23.8	0	23.9	37.3	56.1	100
Zone V (storm surge) (pct)	1,004	2.4	9.8	0	0	0	0	99.8
Zone A (100-yr) (pct)	1,004	39.4	22.5	0	23.0	35.6	51.7	100
Flood insurance claims ^e ('000\$/claim)	188,368	31.3	71.1	0.000	3.3	10.5	32.9	9,139.5
Flood insurance claims, 1991-1995 ('000\$/claim)	29,599	27.4	56.3	0.000	3.1	11.4	31.2	1,845.9
1991-1995 flood claims ('000\$/year/watershed)	1,004	241.9	1,508.3	0	0	0.7	14.3	20,640.9
1991-1995 flood claims, nonzero ('000\$/year/watershed)	635	382.5	1,882.9	0.000	1.1	6.0	49.0	20,640.9
Flood insurance claims, 2016-2020 ('000\$/claim)	41,348	50.0	79.8	0.004	7.3	23.4	64.4	3,000
2016-2020 flood claims ('000\$/year/watershed)	1,004	337.4	1,411.2	0	0.1	6.2	75.7	24,121.5
2016-2020 flood claims, nonzero ('000\$/year/watershed)	785	431.5	1,583.3	0.000	2.0	19.2	135.5	24,121.5
Flood insured value, 1991-1995 ('000'000\$/watershed)	1,004	7.2	37.0	0	0.01	0.1	1.2	757.6
Flood insured value, 2016-2020 ('000'000\$/watershed)	1,004	147.3	438.5	0	0.8	8.5	77.6	5,070.1

Descriptive statistics for Florida, 1995–2020.

^dBank reserves defined as $1 - (\text{total number of offsets sold})/(\text{total number of offsets released})$.

^eAll flood insurance claims from 1985–2020.

TABLE 2. WATERSHED-LEVEL DIFFERENCES BY OFFSET TRADE STATUS

	N	avg	sd	q0	q25	q50	q75	q100
Wetlands developed (acres)								
With wetland bank ^a	96	54.0	72.7	0	4.8	17.9	70.8	239.1
With high development ^b	179	859.5	770.7	253.7	390.3	580.7	958.7	4,812.2
$\mathbb{P}(\text{develop} \text{wet}) \times 100$								
With wetland bank	96	0.8	1.1	0	0.1	0.2	1.0	4.9
With high development	179	15.1	10.1	0.4	7.0	14.2	20.8	57.0
<u>Initial Conditions^c</u>								
Initial wetlands ('000 acres/watershed)								
With wetland bank	96	23.1	76.1	1.4	7.6	10.5	15.9	542.3
With high development	179	10.1	13.8	0.7	3.5	7.0	11.2	140.6
Initial wetlands (pct/watershed)								
With wetland bank	96	44.5	19.4	4.9	29.9	43.3	58.1	95.7
With high development	179	31.0	18.4	5.3	17.8	28.4	39.4	93.5
Initial developed land ('000 acres)								
With wetland bank	96	1.7	2.0	0.03	0.3	0.9	2.1	9.9
With high development	179	10.5	8.8	0.2	5.1	7.7	13.9	49.7
Initial developed land (pct)								
With wetland bank	96	5.7	6.6	0.2	1.2	3.0	8.0	33.1
With high development	179	37.6	23.0	1.5	19.1	32.5	55.7	87.5
<u>Land Ownership</u>								
Watersheds (pct public)								
With wetland bank	96	15.4	21.6	0	0.01	4.5	24.8	95.7
With high development	179	7.1	11.5	0	0.6	2.0	9.7	73.1
Initial public wetlands ('000 acres)								
With wetland bank	96	13.7	74.1	0	0	0.6	3.9	528.1
With high development	179	2.0	6.3	0	0.1	0.4	1.6	68.7
Initial private wetlands ('000 acres)								
With wetland bank	96	9.4	5.5	1.0	5.8	8.4	11.2	31.6
With high development	179	8.1	8.8	0.7	3.3	6.0	9.6	71.9
<u>Flood Risks</u>								
Flood zone (pct/watershed)								
With wetland bank	96	45.8	20.9	0	30.3	46.6	56.7	100.0
With high development	179	33.8	19.8	0.5	20.7	30.6	43.9	99.4
Pre-1996 flood claims ^d ('000\$/yr)								
With wetland bank	96	314.6	2,345.2	0	0	0.1	4.7	22,624.2
With high development	179	412.7	1,552.4	0	0.8	10.7	99.4	13,660.8
Pre-1996 flood insurance (MM\$)								
With wetland bank	96	10.1	77.5	0	0.01	0.1	0.5	757.6
With high development	179	18.8	51.1	0.001	0.6	2.2	9.9	390.6
Post-2015 flood claims ^e ('000\$/yr)								
With wetland bank	96	161.5	541.2	0	0.7	7.1	49.6	3,866.6
With high development	179	798.2	2,445.0	0	24.1	97.3	361.1	24,184.6
Post-2015 flood insurance (MM\$)								
With wetland bank	96	66.0	251.5	0.01	3.7	13.5	32.3	2,267.9
With high development	179	496.4	752.7	3.2	88.2	212.8	522.1	5,070.1

Watershed-level comparison between wetland bank locations and wetland development.
See Table 1 for all data.

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

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

^cInitial measures correspond to 1996 values.

^dPre-1996 flood insurance claims and coverage in 2020 USD, calculated over 1991–1995.

^ePost-2015 flood insurance claims and coverage in 2020 USD, calculated from 2016–2020.

TABLE 3. 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
Average expected parcel utility ('000 USD)	2.2	0.6	0.8	1.6	1.8	1.5	1.9
Std dev expected utility ('000 USD)	7.3	2	2.8	7.5	7.8	7.7	8.3
q50 expected utility ('000 USD)	0.13	0.03	0.04	0.04	0.05	0.02	0.07
q90 expected value ('000 USD)	5.1	1.3	1.9	2.2	2.4	1.3	3.1
q99 expected utility ('000 USD)	35.2	8.4	13.5	44.1	46	50	47.2
Aggregate consumer surplus (bn USD)	4.11	1.12	1.67	2.62	2.34	2.37	2.64
Controls							
Period fixed effects	✓	✓	✓	✓	✓	✓	✓
Water management district fixed effects	✓	✓	✓	✓	✓	✓	✓
Baseline flood risk controls ^b	✓	✓	✓	✓	✓	✓	✓
Lagged development density controls ^c	✓	✓	✓	✓	✓	✓	✓
Lagged demographics ^d			✓	✓	✓	✓	✓
HUC8 fixed effects ^e				✓	✓	✓	✓
Instruments							
Historical sunk capacity		✓	✓	✓			✓
Hausman cost shifters					✓		✓
Government conservation land purchases						✓	✓
First-stage F -stat			115.8	117.3	49.8	8.3	21.3
Observations	758	758	758	758	629	758	629
Adjusted R^2	0.70	0.68	0.70	0.71	0.68	0.64	0.70

Instrumental variable estimates of (14) at the watershed-by-period level for watershed-periods with nonzero development and observed prices. See Section 4.1 for details.

Watersheds correspond to HUC12 units. Periods correspond to intervals in the land cover data (1996–2001, 2001–2006, 2006–2011, and 2011–2016).

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

^bBaseline flood risk controls include percent areas designated as storm surge and 100-year flood zones.

^cDevelopment density controls include percent area developed and the share of developed area that is high development.

^dDemographic controls are population and median income.

^eHydrological unit code (USGS, 2013).

TABLE 4. ESTIMATED WETLAND BANK COSTS

	N	mean	sd	q10	q25	q50	q75	q90
First-stage entry probabilities								
$p_{\{\text{enter}\}}$, firm 1	29	0.15	0.09	0.06	0.08	0.15	0.21	0.26
$p_{\{\text{enter}\}}$, firm 1, duopoly	6	0.18	0.07	0.11	0.16	0.19	0.23	0.25
$p_{\{\text{enter}\}}$, firm 2, duopoly	6	0.11	0.05	0.05	0.08	0.12	0.15	0.16
$p_{\{\text{enter}\}}$, firm 1, oligopoly, at least three firms	15	0.19	0.09	0.10	0.14	0.17	0.21	0.28
$p_{\{\text{enter}\}}$, firm 2, oligopoly, at least three firms	16	0.18	0.07	0.10	0.12	0.17	0.23	0.27
$p_{\{\text{enter}\}}$, firm 3+	55	0.17	0.06	0.08	0.12	0.19	0.22	0.25
Value functions								
$\mathbb{E}[V]$, firm 1	29	15.96	22.01	1.71	4.23	9.00	18.54	34.87
$\mathbb{E}[V]$, firm 1, duopoly	6	29.13	39.59	9.28	10.58	14.05	17.53	64.08
$\mathbb{E}[V]$, firm 2, duopoly	6	5.84	3.99	2.11	3.31	4.72	8.74	10.69
$\mathbb{E}[V]$, firm 1, oligopoly, at least three firms	15	15.05	15.38	1.86	3.42	9.00	25.13	36.22
$\mathbb{E}[V]$, firm 2, oligopoly, at least three firms	16	16.67	17.31	0.95	3.59	8.38	28.52	39.12
$\mathbb{E}[V]$, firm 3	55	9.13	22.30	0.43	1.10	2.84	8.36	14.57
Parameter estimates								
$\mu_{\kappa}(s_{mt})$, firm 1	29	24.57	6.23	21.74	22.97	22.97	24.80	27.96
$\sigma_{\kappa}(s_{mt})$, firm 1	29	8.36	0	8.36	8.36	8.36	8.36	8.36
$\mu_{\kappa}(s_{mt})$, firm 1, duopoly	6	26.61	6.94	22.35	23.44	24.80	24.80	32.67
$\mu_{\kappa}(s_{mt})$, firm 2, duopoly	6	20.79	6.94	16.54	17.60	18.97	18.98	26.86
sig, duopoly	12	7.83	0.54	7.31	7.31	7.83	8.36	8.36
$\mu_{\kappa}(s_{mt})$, firm 1, oligopoly, at least three firms	15	26.00	6.01	22.25	22.97	24.78	24.80	34.25
$\mu_{\kappa}(s_{mt})$, firm 2, oligopoly, at least three firms	16	21.09	6.85	16.54	17.15	18.97	18.98	34.72
$\mu_{\kappa}(s_{mt})$, firm 3	55	12.15	7.44	6.52	8.35	8.35	24.10	24.10
$\sigma_{\kappa}(s_{mt})$, oligopoly, at least three firms	86	4.81	2.29	3.12	3.12	3.12	7.31	8.36
Implied costs								
Realized entry cost estimate (MM/bank)	99	6.62	8.96	-3.53	0.40	2.51	8.66	40.55
Est entry costs per credit ('000/bank)	99	21.51	37.20	-67.59	1.43	13.48	31.91	176.56
Comparison with contract data								
Observed entry costs (MM/bank)	79	5.29	6.09	0.26	1.42	2.86	7.18	36.16
Observed entry costs per credit ('000/bank/credit)	79	23.95	23.27	1.76	9.20	15.99	31.17	116.67
Observed construction costs (MM/bank)	86	1.61	2.50	0.04	0.36	0.97	1.81	16.16
Observed land costs (MM/bank)	95	5.05	10.53	0.02	0.57	1.89	5.53	89.19
Implied markup	79	2.97	6.18	1.06	1.44	1.65	2.70	55.50
Rate of return on capital (pct)	79	7.78	7.18	0.58	3.71	5.14	10.44	49.43
Rate of return on capital (pct), firm 1	24	8.18	9.48	1.41	3.77	4.62	10.43	49.43
Rate of return on capital (pct), firm 2	18	8.05	5.71	1.53	4.24	7.01	10.99	22.32

First and second-step cost estimates for wetland banks. See Section 4.2 for details.

TABLE 5. ESTIMATED LOCAL FLOOD DAMAGE FUNCTIONS

	(1)	(2)	(3)	(4)	(5)	(6)
Development on wetlands (ζ_d)	0.484 (0.123)	0.267 (0.109)	0.253 (0.111)	0.253 (0.120)	0.209 (0.105)	0.199 (0.133)
Wetland bank area (ζ_b)	0.121 (0.044)	0.046 (0.039)	0.046 (0.038)	0.021 (0.039)	-0.021 (0.054)	0.017 (0.041)
Baseline flood claims (1991-95)		0.439 (0.028)	0.419 (0.029)	0.421 (0.032)	1	0.376 (0.034)
Identifying assumption	OLS	AR(1)	AR(1)	AR(1)	LD	AR(1)
Controls						
Water district fixed effects	✓	✓	✓	✓	✓	✓
Baseline flood risk	✓	✓	✓	✓	✓	✓
Baseline dev density	✓	✓	✓	✓	✓	✓
Other development	✓	✓	✓	✓		✓
Demographic controls	✓		✓	✓	✓	✓
HUC8 FEs				✓	✓	✓
Estimation sample						
Wetland development	✓	✓	✓	✓	✓	✓
Baseline insurance					✓	✓
Implied damages						
0%	0.8	0.2	0.1	0.1		0.2
10%	28.1	5.5	5.3	3.9		4.1
20%	54.7	9.0	9.7	9.4		11.8
30%	81.3	16.5	17.5	18.4		25.4
40%	119.9	44.4	43.7	41.2		53.9
50%	183.5	140.7	134.0	104.7		127.8
60%	291.4	327.5	325.7	314.5		314.5
70%	543.7	709.7	708.2	798.1		759.1
80%	1,277.8	1,702.7	1,652.7	2,124.1		1,963.5
90%	4,717.6	5,686.4	6,149.1	10,014.2		7,762.0
95%	14,742.6	13,554.2	13,225.6	35,423.1		22,111.5
97.5%	47,563.3	36,065.4	36,071.5	86,793.2		63,299.0
99%	264,588.9	181,535.6	247,786.1	476,471.2		381,874.5
99.9%	6,947,372.0	1,914,000.0	1,157,514.0	2,837,490.0		1,591,938.0
100%	51,110,292.0	10,377,320.0	3,410,034.0	4,581,752.0		2,195,878.0
Observations	1,047	1,054	1,047	1,047	866	896
Adjusted R ²	0.433	0.531	0.532	0.592	0.268	0.552

Estimates of (20) at the local watershed level for watersheds with at least one acre of development on wetlands. The outcome is flood insurance claims after the market (2016–2020) for properties built prior to the market (1995); see Table A4 for other outcomes. Columns (5) and (6) restrict the sample to watersheds with nonzero flood insurance policies in 1995. Implied damages report the distribution of marginal damages (at observed development under the market) per acre wetland developed over all watersheds with at least one acre of wetland developed (1,047 watersheds).

Robust (HC1) standard errors clustered at the HUC12 level in parentheses.

TABLE 6. ESTIMATED MARGINAL DAMAGES

	Per acre	Per credit
0%	0.10	1.16
5%	3.21	23.65
10%	5.33	41.50
15%	7.28	58.08
20%	9.69	82.22
25%	12.11	103.75
30%	17.52	143.24
35%	24.95	220.00
40%	43.70	360.17
45%	72.81	660.43
50%	134.00	1,109.29
55%	225.25	1,669.29
60%	325.74	2,560.92
65%	469.26	3,701.26
70%	708.20	5,734.15
75%	1,022.76	8,628.47
80%	1,652.68	13,015.81
85%	2,660.14	24,929.85
90%	6,149.11	49,513.96
95%	13,225.63	115,548.70
97.5%	36,071.49	319,856.40
99%	247,786.10	1,971,415.00
99.9%	1,157,514.00	10,838,671.00
100%	3,410,034.00	16,469,952.00

Estimated marginal damages across local watersheds in our data with at least one acre of wetland development from 1996–2016 (1,047 watersheds).

Calculated as the expected net-present-discounted marginal damage (excess insured flood damage) of permanently developing an acre of wetland (column 1) or using a wetland offset (column 2). See Section 5.2 for details.

TABLE 7. WELFARE AND OFFSET MARKET DESIGN

	Market	Pigouvian tax	Uniform tax
Wetlands developed (acres)	140,653.5	81,339.9	108,788.4
Wetlands offsets used (credits)	16,705.5	9,369.0	11,991.0
Gains from trade			
Developer values (MM)	2,801.1	1,938.2	2,591.9
Supply costs (MM)	609.7	414.5	480.6
Private gains from trade (MM)	2,191.4	1,523.7	2,111.4
Distributional outcomes			
Consumer surplus (MM)	1,672.2	975.1	1,265.9
Producer surplus (MM)	519.2	278.4	491.7
Tax revenue (MM)	0	270.2	353.7
Externalities			
Flood damage to existing structures (MM)	−1,604.7	−258.0	−1,420.7
below 99.9%-ile	−1,602.4	−258.0	
below 99%-ile	−1,393.9	−258.0	
below 97.5%-ile	−1,012.2	−257.8	
7% discount rate	−1,151.9		
3% discount rate	−2,565.6		
Welfare (MM)	586.7	1,265.7	690.6


Value in millions of 2020 USD.

Market outcomes from 1995–2018 at observed offset prices (column 1), offset prices with local Pigouvian taxes (column 2), and offset prices with a uniform tax (column 3).


Flood damages’ net present discount value calculated using a 5% 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 \$30,000/offset.

A. Florida Offset Markets



B. Example: HUC 03100205



C. Observed development and wetland mitigation banking

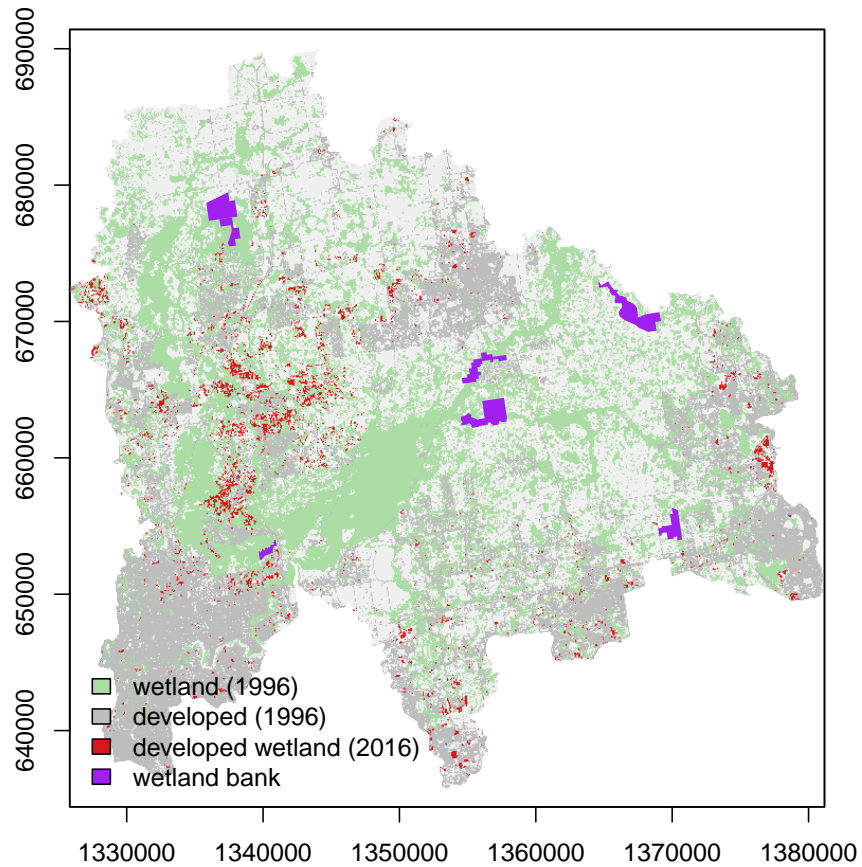


FIGURE 1. LOCATIONS OF WETLAND DEVELOPMENT AND RESTORATION

An example of our data on land use and wetland offsets within an offsets market. Initial wetland (green) and developed (grey) pixels in 1996, new development on wetlands from 1996–2016 (red), and wetland bank parcels established by 2018 (purple).

Online Appendix Figures A4.1–30 replicate this map for every market in our study.

Table 2 reports average differences between all watersheds, watersheds with development (red pixels), and watersheds with wetland mitigation bank sites (purple pixels).

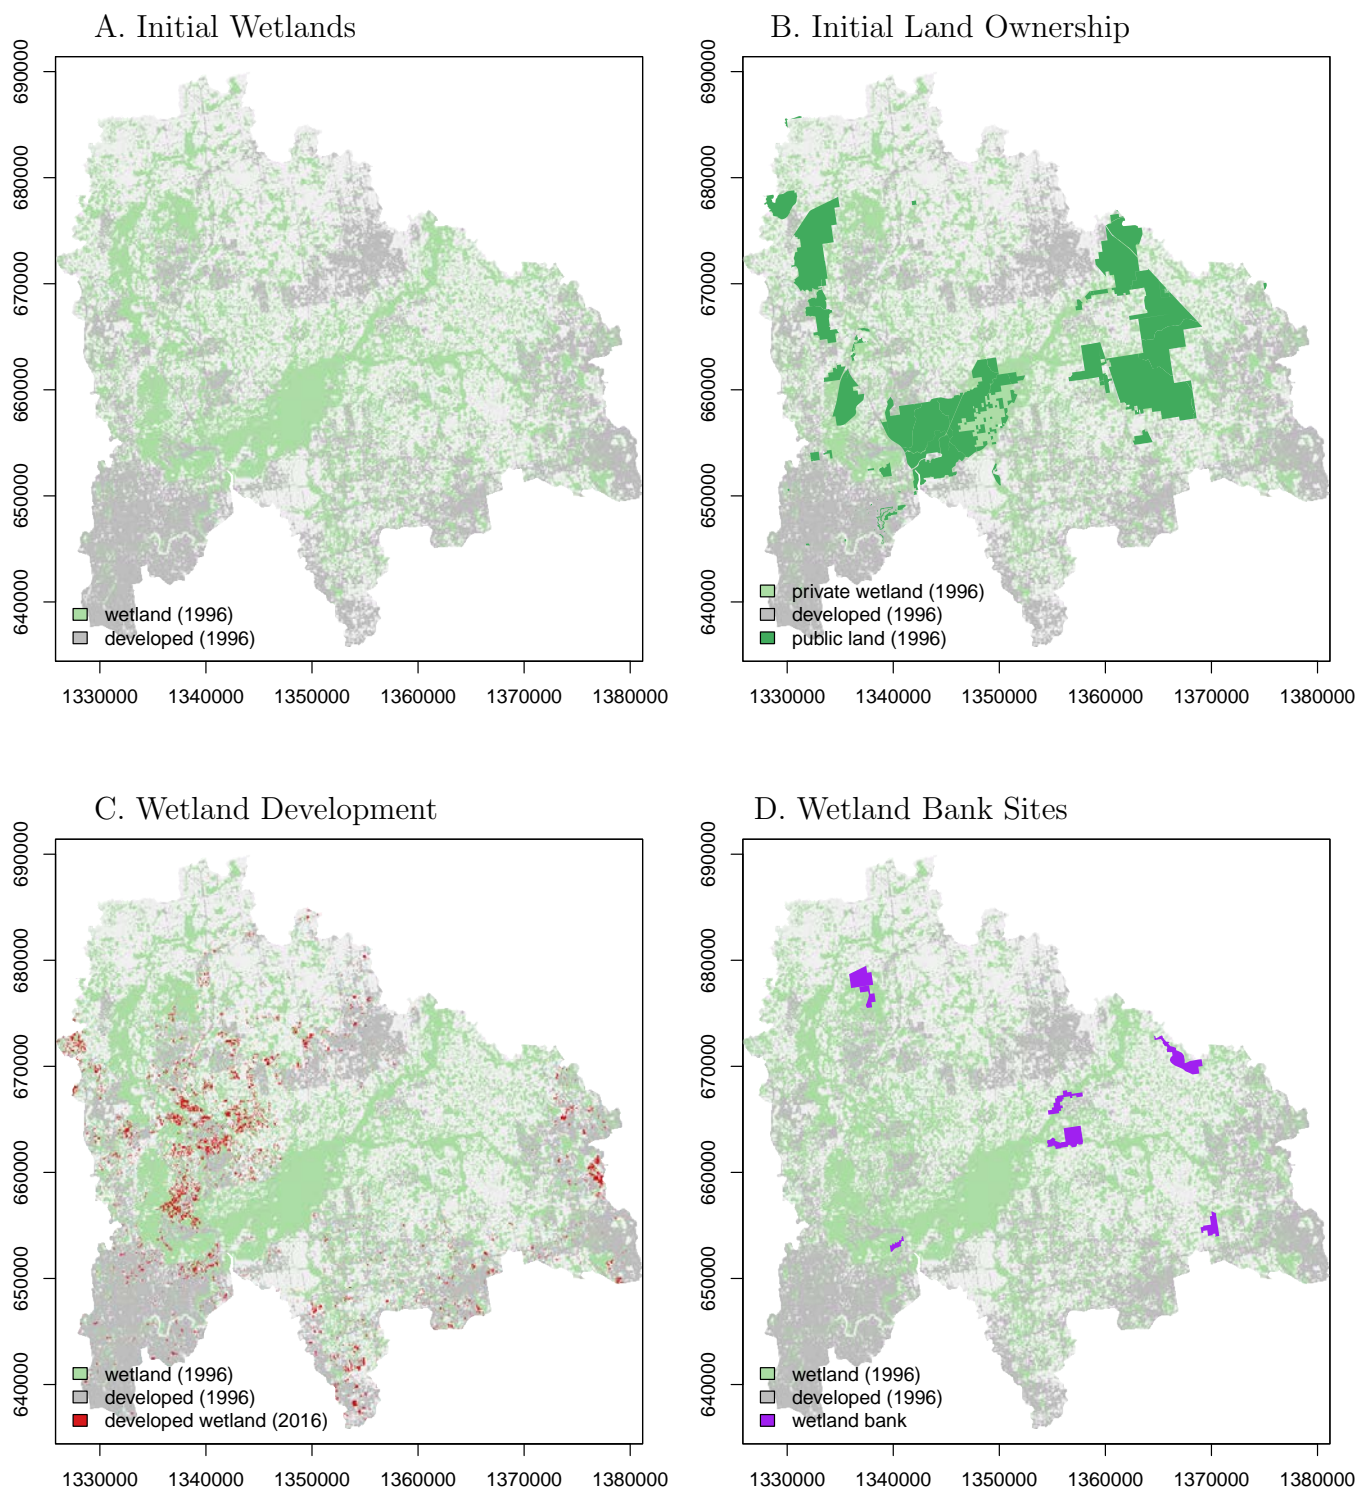
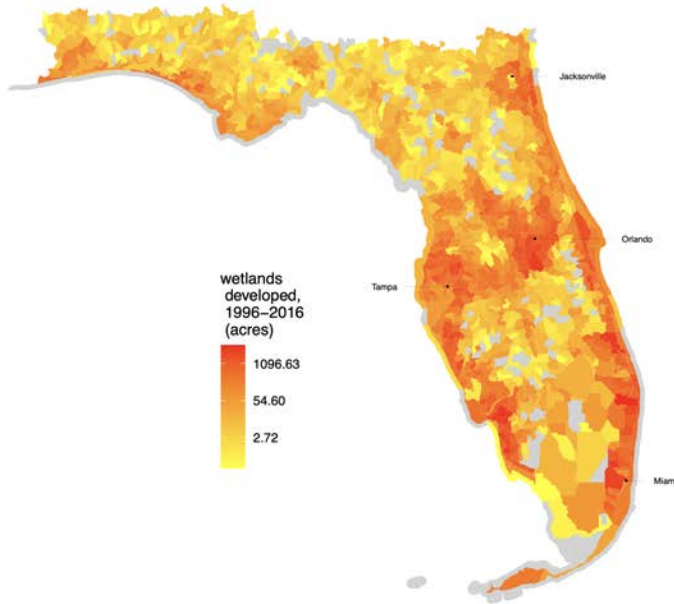


FIGURE 2. INITIAL CONDITIONS, OWNERSHIP, DEVELOPMENT, AND RESTORATION

An example of our data on land use, ownership, and wetland offsets within a market.

- A. Initial wetland (green) and developed (grey) pixels in 1996.
- B. Initial public land (dark green) in 1995.
- C. New development on wetlands from 1996–2016 (red).
- D. Wetland bank parcels established by 2018 (purple).

A. Development on Wetlands



B. Estimated Developer Surplus


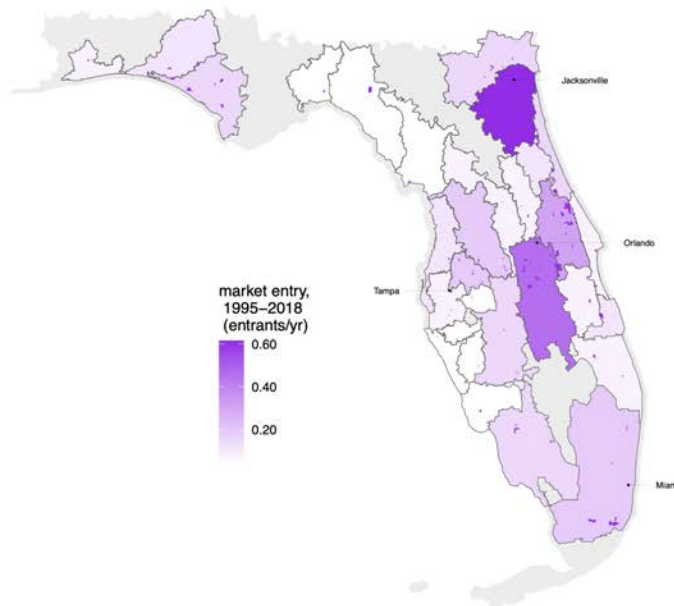


FIGURE 3. DEVELOPMENT ON FLORIDA WETLANDS

A. Map of local watershed development between 1996–2016. Local watersheds colored by decile of $\ln(\text{acres of wetlands developed})$.

B. Estimated private consumer surplus for land developers purchasing offsets and calculated with (21), ordered left to right by trades' decreasing estimated surplus, 1995–2018.

A. Entry Probabilities



B. Estimated Producer Surplus and Costs




FIGURE 4. WETLAND MITIGATION BANKS

A. Map of estimated local watershed entry probabilities colored by decile (right).

B. Estimated per-credit costs and transaction values for wetland banks, calculated with (A2) and (A3) and ordered left to right by increasing simulated price per credit.




FIGURE 5. REALIZED PRIVATE GAINS FROM TRADE

Land developers' private values, transaction prices, and wetland banks' private costs.

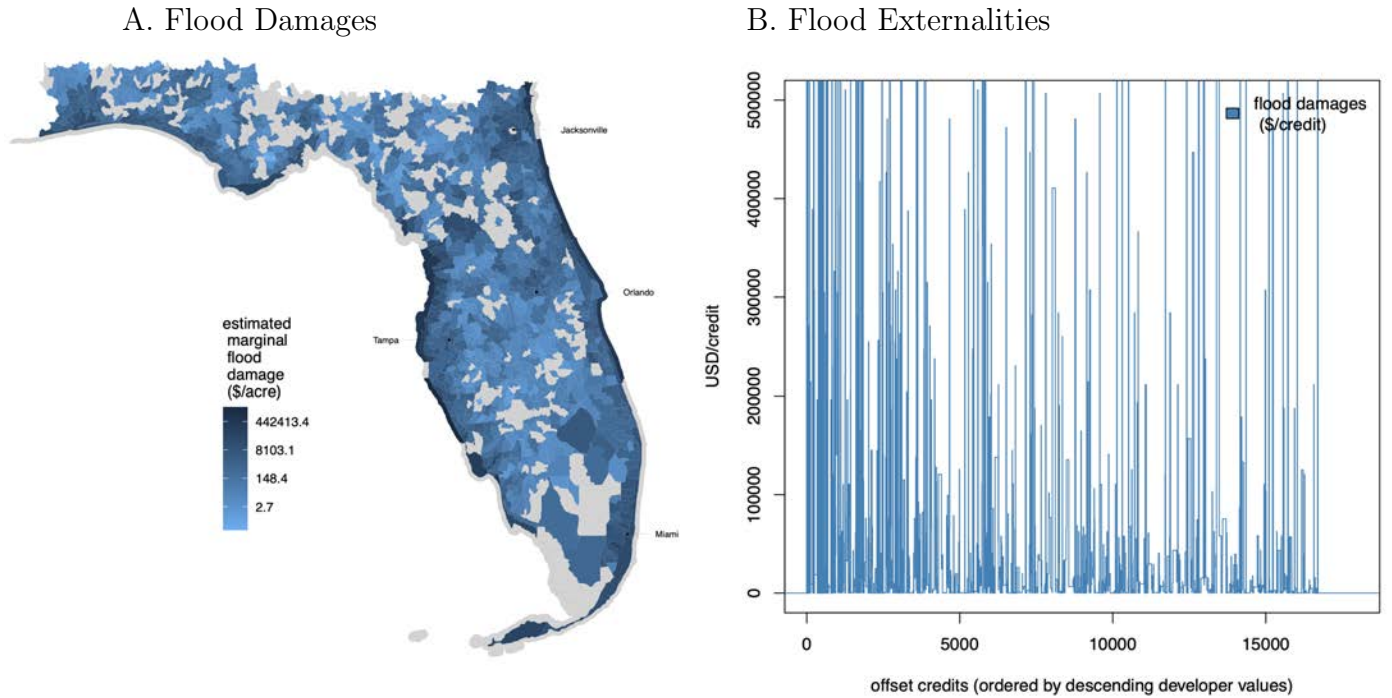


FIGURE 6. REALIZED FLOOD DAMAGES

A. Map of estimated marginal flood damages at the watershed level for wetlands with nonzero (at least one acre) of wetlands developed within an offset market.

B. Estimated average flood damage for each wetland under the market from 1996–2016, calculated with (18) and sorted by descending private surplus from Figure 3, Panel B.

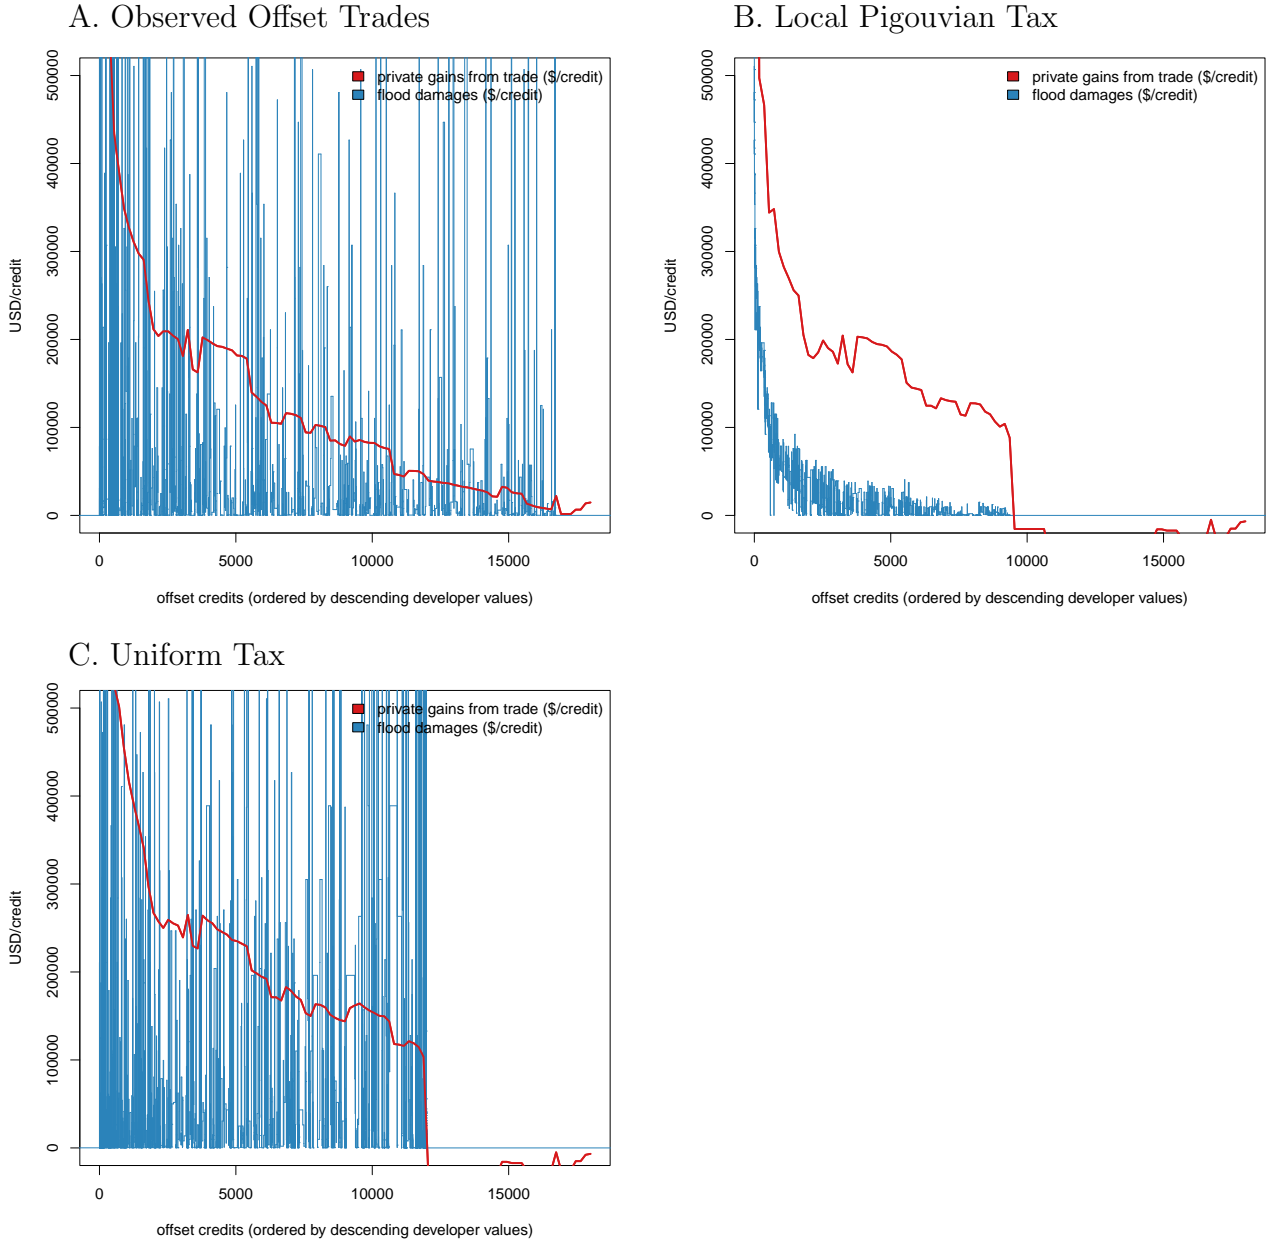


FIGURE 7. PIGOUVIAN REDESIGN

A. Estimated flood damages from Figure 6, Panel B, plotted against the private gains from trade (i.e., the difference between the developer values and bank costs in Figure 5).

B. Estimated private gains from trade and flood damages under the Pigouvian flood protection taxes at the local watershed level, sorted by descending developer value.

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 developer value.

See Section 5.3 for details.