

Economic Tradeoffs and Institutional Guidance in Public Resource Allocation: Evidence from U.S. Public Lands

Maya A. Norman*

Job Market Paper

[Click here for the most recent version](#)

October 15, 2026

Abstract

The US federal government owns nearly half of the Western United States, one of the least developed contiguous landscapes in the world. Public agencies, rather than markets, allocate these lands between development and conservation. The Bureau of Land Management (BLM) is the primary agency making this tradeoff, deciding where mineral extraction that generates billions in public revenue should proceed and where it should be forgone for conservation. Despite the importance of the BLM's allocation decisions, the mechanisms underlying them have never been studied. Federal law requires the BLM to maximize the land's market and non-market resource values and publicly justify its choices to stakeholders. But the law does not specify how to value non-market resources or reconcile competing demands from pro-conservation and pro-industry stakeholders. As a result, roughly 150 local offices retain wide discretion over these judgments. Until now, systematically studying how offices exercise that discretion has been infeasible because the evidence is scattered across sprawling planning records and maps. I develop a framework using AI agents to transform these records into a form suitable for discrete-choice econometric methods. I then apply this framework to 11 local offices and roughly 90,000 pages of planning records, recovering how much revenue each office would forgo to (i) conserve non-market resources and (ii) follow established conservation standards rather than justify departing from them to stakeholders. I find that the median office is willing to forgo roughly \$60,000 in annual revenue to conserve one acre. I also find that in seven out of ten decisions, the median office passes up the resource-value-maximizing choice to remain closer to an established standard. This tendency to follow standard practice matters for policy. I compare two ways BLM headquarters can steer local offices. One is to recommend a specific conservation practice, such as a one-mile no-drilling buffer around birds' nests. The other is to use a price-like instrument that specifies the value of conserving a resource. Counterfactuals show that the recommended practice can be as cost-effective as the price-like instrument in achieving the same conservation gains.

Acknowledgments: I thank Bernard Salanié, Pietro Tebaldi, Jeffrey Shrader, Wolfram Schlenker, Mike Ting, Pierre Bodéré, Haoge Chang, Gautam Gowrisankaran, Simon Greenhill, Wojciech Kopczuk, Marguerite Obolensky, Sebastián Otero, Brendan O'Flaherty, Clara Berestycki, Abhishek Deshwal, Matthew Goroff, Dylan Hogan, Roberto Zúñiga and participants of the following Columbia University Colloquiums: IO and Sustainable Development, for very helpful feedback.

*Sustainable Development, Columbia University, email: man2185@columbia.edu

1 Introduction

The U.S. federal government owns nearly half of the Western United States (345 million acres) (Hanson & Vincent 2020). Public agencies, rather than markets, allocate these lands among competing uses. The Bureau of Land Management (BLM) is the primary public land manager in the Western US. It determines where mineral development generating billions in public revenue annually can proceed and where development is restricted to conserve one of the world’s largest intact landscapes (Office of Natural Resources Revenue 2025, Doherty et al. 2024). Despite the importance of these decisions, economists know little about what drives them (Libecap 2026).

Federal regulations require the BLM to maximize market and non-market resource values. But they do not specify what value to place on non-market resources, such as a bald eagle’s nest. As a result, roughly 150 local offices have considerable discretion over how to protect these resources. An office might follow a state wildlife agency’s recommendation to restrict drilling within a quarter mile of the nest, or choose a different buffer. Whatever buffer it chooses, regulation also requires the office to publicly justify its decision to stakeholders, including oil and gas companies and conservation groups.¹ Research on public bureaucracies suggests that this accountability requirement can make offices care not only about maximizing resource values, but also about how easily they can defend their choices (Prendergast 2003, Dixit 2002). The office may therefore follow the quarter-mile recommendation even if local benefits and costs would favor a different restriction. This is because citing the recommendation is easier than defending its own assessment of the benefits and costs. I study how offices navigate these institutional demands by estimating how their choices respond to local benefits and costs and whether the ease of defending those choices also matters.

Studying these decisions requires observing what offices chose, what it could have chosen, and the expected consequences of each alternative. Assembling this information is difficult because land-use allocation decisions involve many dimensions. By default, BLM lands are available for development. Allocation decisions therefore amount to deciding where to restrict development. To protect non-market resources (bald eagles, prairie dogs, caves, scenic landscapes, and others), BLM offices choose which market activities (oil, gas, coal, wind, solar, grazing, and others) to restrict and how far those restrictions extend around each resource. Federal regulations require the BLM to document its chosen allocations, the alternatives it considers, and the expected consequences of each for market and non-market resources. But this information is scattered across thousands of pages of text and maps in each office’s planning record.

I use large language models (LLMs) to transform these planning records into data suitable for discrete-choice econometric methods. AI agents read 90,000 pages of planning documents in three stages. They first read the full record to locate relevant passages. They then reread and label those passages according to what each establishes. Finally, agents use the second-stage labels to select and reread the passages needed to construct the estimation sample. From these passages, they recover the chosen and considered allocations, along with the BLM’s projections of development, revenue, and disturbance under different land-use scenarios. After this stage, I verify every coded value ultimately used to construct the estimation sample against the original planning records. I then map the geographic extent of restrictions under each allocation the BLM chose or considered. I combine these maps with the BLM’s development projections to measure the revenue forgone and disturbance avoided under each alternative. The resulting dataset contains 12,509

¹See appendix A and D.1.

choice situations across 11 local offices in Colorado, Idaho, New Mexico, and Wyoming. For each choice situation, it identifies the chosen allocation and the feasible alternatives. For every alternative, it reports the expected development revenue forgone (the option’s cost) and the expected disturbance avoided (the conservation it buys).

I use these data to estimate a discrete-choice model that identifies offices’ (revealed) preference for (i) conservation and (ii) following established conservation standards. I find substantial variation in how much development revenue offices are willing to forgo for conservation, ranging from \$12k to \$1.6 million annually per acre. This variation occurs both across types of non-market resources within an office and across offices for the same non-market resource type. Established standards also have a substantial influence on choices. At the median office, the fitted model predicts that in seven out of ten decisions, the office selects an allocation closer to an established standard rather than the allocation that maximizes conservation value net of forgone revenue.

This preference to follow standard practice matters for policy. I compare two ways BLM headquarters can steer local offices. The first is a quantity-like instrument that recommends a specific conservation practice, such as a one-mile no-drilling buffer around birds’ nests, without requiring offices to adopt it. The second is a price-like instrument that sets a common value for bird conservation across offices. I set the common value so that both instruments achieve the same aggregate conservation gains. I find that the price-like instrument is no more cost-effective than the quantity-like instrument in increasing the number of baby birds protected. But the instruments differ in where they achieve gains. Half of the quantity-like instrument’s gains occur in Buffalo, WY, whereas half of the price-like instrument’s gains occur in Casper, WY and none in Buffalo, WY. This stark contrast reflects differences in how offices value bird conservation. Buffalo’s valuation of bird conservation exceeds the common value, so the price-like instrument does not induce additional protection there. Raising the common value enough to induce more conservation in Buffalo would also increase conservation elsewhere, raising costs and exceeding the aggregate conservation goal. The quantity-like instrument, however, can change Buffalo’s choices by giving it a new standard to follow, without requiring headquarters to raise the valuation applied to every office.

Relevant Literature There is no empirical research on how the Bureau of Land Management makes land-use decisions, even though these choices determine the allocation of vast public resources across competing users and generations (Libecap 2026). This absence mirrors a broader gap in empirical studies of the regulatory process despite longstanding recognition that such work provides “important, useful insights into what is actually happening” (Joskow & Noll 1981, Magat et al. 1986, Oates & Portney 2003). In the empirical literature on environmental regulatory behavior only two papers—Magat et al. (1986) and Cropper et al. (1992)—study how environmental agencies make decisions in the US. These papers apply revealed preference techniques pioneered by McFadden (1976, 1975) to examine how agency discretion is applied in actual decision making. McFadden developed those techniques for a public bureaucracy choosing among discrete, spatially explicit alternatives, recovering the benefit–cost calculus behind freeway route selection by the California Division of Highways—testing which benefits and costs entered the agency’s decision, and whether political factors moved it. The zoning problem studied here has that structure, and Section 2 sets it out. Magat et al. (1986) develop a revealed-preference approach to recover how regulatory agencies make rules. They apply their method to EPA rulemaking decisions to infer how EPA regulators trade off among

the various factors, including information and input supplied by regulated industries. Cropper et al. (1992) use a revealed-preference approach to ask to what extent the EPA weighs a pesticide’s economic benefits against its health risks when deciding whether to ban it. This analysis assesses the extent to which the EPA meets their congressional mandate when regulating pesticides. Along a similar vein, this paper uses zoning decisions to assess the extent to which the economic benefits of development are weighed against the costs of development in terms of natural capital loss.

An existing literature studies tradeoffs between conservation and alternative economic uses of land. A central concern in this literature is additionality: whether conservation policies protect land that would otherwise have been converted to economic use. Conservation creates real opportunity costs when it displaces development (Barbier & Burgess 2001), motivating policies that compensate landowners or governments for forgone development opportunities (Angelsen et al. 2012). Economists have studied whether conservation investments actually deliver environmental benefits (Ferraro & Pattanayak 2006), how selection into conservation contracts can limit additionality (Aspelund & Russo forthcoming), and why protected areas may fail to prevent land conversion when protections are weak or targeted toward low-threat locations (Reynaert et al. 2024). Evidence from federal biodiversity regulation similarly suggests that conservation boundaries may be drawn in part to avoid large economic costs (Frank et al. 2025). This paper extends the additionality concern to discretionary public-land planning. Rather than asking whether conservation programs change private land-use behavior, I ask whether public-land managers choose conservation where doing so requires forgoing meaningful development value.

This question connects to a growing literature on decentralized public decision-making in settings where local governing bodies shape investment and land-use outcomes. While this literature has longstanding theoretical roots (Oates 1972, Fischel 2001), recent empirical work increasingly uses observed public decisions to recover the objective functions and constraints of local decision makers (Bodéré et al. 2026, Kashner 2026, Slattery 2025, Parkhomenko 2023). For example, Bodéré et al. (2026) estimate a dynamic model of drinking-water infrastructure investment in which locally elected officials trade off residents’ welfare gains from improved water quality against administrative costs and electoral incentives. Similarly, Kashner (2026) studies local government approval of wind farms, modeling projects as built only when they are both profitable for developers and permitted by local governments that weigh local externalities against developer payments. Slattery (2025) models subsidy competition for firm location as an auction and uses observed subsidy deals to recover local governments’ willingness to pay for different firms. Parkhomenko (2023) develops a quantitative spatial equilibrium model in which land-use regulation is an endogenous policy choice. In these papers, public decisions are endogenous to local incentives and constraints, providing modern examples of structural and revealed-preference approaches to decentralized public decision-making in environmental, infrastructure, tax, and land-use settings. My paper shares this focus on public decisions as endogenous determinants of outcomes, but differs by focusing specifically on the administrative planning process. I treat the planning record itself as data, using it to reconstruct the feasible choices, information environment, regulatory constraints, and perceived payoffs that structure public-land managers’ decisions. Rather than evaluating these decisions using only *ex post* estimates of benefits and costs, I focus on the benefits and costs documented at the time decisions were made. Because *ex ante* assessments may differ substantially from realized outcomes, this approach provides a faithful representation of the decision environment facing public-land managers and allows me to estimate the tradeoffs implied by the choices

themselves.² I therefore study planning as an object of economic analysis: given the information available and policy instruments at hand, what conservation-development tradeoffs did public-land managers make in their zoning decisions?

The second force in the model, the cost of justifying a choice, connects this paper to research on officials whose decisions can be challenged. McCubbins & Schwartz (1984) and McCubbins et al. (1987) argue that elected officials control agencies through procedure: rules that let citizens and interest groups examine an agency’s decisions and seek remedies do the monitoring that legislators do not do themselves. The comment process a field office must go through is a procedure of this kind. Officials who expect to be challenged choose with the challenge in mind. In Leaver (2009), bureaucrats who want to avoid criticism take the decisions that keep interest groups quiet. In Prendergast (2003), bureaucrats who are overseen through complaints have a reason to give the complaining party what it asks for. A field office is criticized from both sides, so no buffer keeps every commenter quiet; an established standard instead supplies an answer the office can give to both. A second line of work explains why decision makers stay with established practice. Scharfstein & Stein (1990) show that managers who care about their reputations copy the investments of other managers and ignore what they privately know. Bernheim (1994) shows that people with different tastes conform to one standard of behavior when others judge them by their actions. Frakes (2013) finds that practice follows the standard a decision is judged against: when states began to hold physicians to national rather than local standards of care in malpractice law, treatment rates in those states moved toward the national rates. These papers establish the mechanism in theory or detect it in outcomes. Here the planning record shows offices citing standards to answer their critics, and I estimate what each office is willing to give up in development revenue to stay near its target, alongside its valuation of the natural capital itself. An office with a high cost of departure holds to its target regardless of what is at stake at the feature; an office with a low cost weighs those stakes; an office in between does some of both. A high cost may reflect how hard it is to defend a buffer of the office’s own or how much time and staff it takes to assess a specific buffer around each feature; either way, a high cost means the office generally holds to its target and does not weigh costs and benefits feature by feature.

Closer to this paper’s specific context, the limited empirical literature on federal land management largely examines how management decisions—such as permitting grazing, requiring environmental review, or implementing wildfire fuel treatments—affect downstream outcomes, including drilling activity, wildfire risk, and ecological change, while taking the land-management decision-making process itself as given (Lewis 2019, Edwards et al. 2019, Walls et al. 2020, Jakus & Akhundjanov 2019, Rimbey et al. 2007, Egan & Watts 1998, Jones & Carter 2016, Davies et al. 2014, Steinbach & Thomas 2007, Cochrane et al. 2012). Because outcomes arise only where an activity is permitted, the effects documented in this outcomes-oriented literature are downstream of the BLM’s initial zoning decisions. For example, permitting reforms can only affect oil and gas activity on parcels zoned as available for drilling, and subsidies for renewable energy can only influence development on lands the BLM has already designated as eligible for renewables. My analysis provides an upstream complementary perspective by identifying how planners decide which parcels

²Rollet (2025) makes a related point in the context of urban redevelopment, arguing that New York City’s zoning code may have been less costly, and potentially beneficial, under the economic and industrial conditions prevailing when it was adopted. The policy became increasingly costly as demand, land values, and the externalities from different land uses changed. Because zoning adjusts slowly, observed outcomes today can reflect institutional persistence rather than contemporaneous policymaker preferences. In my setting, using realized outcomes alone to infer public-land managers’ objective function could similarly conflate the tradeoffs embedded in the original planning decision with later constraints, shocks, and implementation dynamics.

become eligible for use. In doing so, it provides a critical component for understanding how regulatory or technological changes would alter outcomes (Joskow & Noll 1981). For instance, a technology that raises the revenue expected from a well while reducing the disturbance it causes raises the cost of protecting a feature and lowers the benefit. An office that weighs the stakes then reopens features whose forgone revenue now exceeds the value it places on the disturbance avoided, and the technology's environmental gain is undone by the rezoning. An office that instead follows standard practice, zoning a little less protection than its peers whatever the stakes, leaves its plan alone, and its jurisdiction realizes the gain.

This paper also contributes to a debate about the contribution of public participation to agency decision-making. While canonical political-economy models posit that stakeholder involvement in regulatory processes can be distortionary (Becker 1983, Peltzman 1976, Stigler 1971, Woods 2013), qualitative work in environmental management suggests that incorporating local knowledge may, in practice, generate substantial efficiency gains (Martin et al. 1996, Oates & Portney 2003). Rather than taking a stand on whether stakeholder involvement is distortionary or efficiency-enhancing, my positive revealed-preference approach allows the data to speak. I recover the implicit decision rules produced by the existing regulatory process, doing so in a way that accommodates any underlying political equilibrium (Becker 1983, Peltzman 1976, Stigler 1971). The recovered rules provide a basis for evaluating, *ex post*, whether the process generates undesirable or systematically biased allocations, while avoiding any presumption *ex ante* about the quality or desirability of the regulatory process. In addition, this approach does not require *ex ante* assumptions about the identity of stakeholders and thus can be readily applied to any land-use plan without requiring supplemental data on stakeholder participation—a challenge documented in the land-management literature (Martin et al. 1996, Oates & Portney 2003).

2 The Model


Field offices decide if and how to protect non-market resources in their jurisdiction, weighing agency precedent against the benefits and costs of protection decisions—the non-market resources protected and the revenue forgone. Four Rivers, Idaho, bars surface disturbance around a burrowing owl nest during the nesting season within a quarter mile of the nest. Royal Gorge, Colorado, imposes the same restriction within 300 feet. This section develops a model that separates two forces behind differences like this one: the pull of the standard an office follows, regardless of the stakes; and its response to the net benefits at a particular feature—how much conservation value a wider buffer protects and how much more revenue it gives up. The model asks how these two forces combine to produce the conservation decisions we observe.

2.1 The Decision Problem

A field office’s central allocation decision is where to withhold land from development, and it makes that decision *natural feature* by *natural feature*—raptor nest by stream by steep hillside. For each feature f , the office chooses two things: a *regulatory regime* r , which determines whether and how strictly development is restricted, and a *buffer* b , which determines how far that restriction extends around the feature. Diagram 1, at right, illustrates the anatomy of such a decision.

I cast the allocation decision for each feature as a discrete choice in the framework of McFadden (1974): an agent selects one alternative from a feasible set, the econometrician observes both the choice and its attributes—including the costs and benefits—of every alternative available, and preferences follow from what the agent chose over other feasible options (Train 2009). I also let an alternative’s distance from the standard the office follows enter the office’s payoff, capturing a preference for conformity (Bernheim 1994); I measure that standard with a proxy, the buffer the other offices chose, plus a correction for the standard the office actually follows. The central object I recover is the exchange rate between the stakes of protection: the revenue each field office is willing to forgo to protect natural capital.

Diagram 1: A Zoning Decision



Notes: The blue shaded disk is the policy footprint—the area restricted by the regulatory regime r and defined by the buffer b . The dashed rings are buffers the office could have chosen.


2.2 Supporting Facts

The decision problem of Section 2.1 is an abstraction, and its elements are modeling choices. Four facts in the data support them. First, offices protect less as the revenue at stake rises, though not all at the same rate, motivating office-specific exchange rates. Second, zoning is organized feature by feature, motivating the feature as the choice situation. Third, every considered configuration pairs a regime with an extent anchored on the feature, motivating regime–buffer pairs as the alternatives. Fourth, choices cluster at a handful of protection extents shared across offices, often with more than one common extent for the same kind of feature. This motivates modeling a cost of departing from standard practice while allowing offices to follow different standards.

The preferences: office-specific exchange rates. Figure 1 plots the share of each township³ closed to fluid leasable development against each township’s development potential. Eight of the eleven offices in the estimation sample protect a steadily smaller share as projected revenue rises; at Buffalo, Wyoming, and at Colorado River Valley and Uncompahgre, Colorado, the share is flat or rises. Panel (a) shows the contrast within one state: Lander and Rock Springs protect less as revenue rises, while Buffalo protects more. Pooling all eleven offices, the share closed falls from 41% at about \$260 a year to 18% at about \$18 million. Two aspects of the figure bear directly on the exchange rate I seek to recover. *First*, the cost of conservation matters: as the revenue at stake rises, most offices protect less. A model that omitted forgone revenue would therefore miss a central decision-making margin. *Second*, the tradeoff differs across offices. In panel (b), the office curves fall at visibly different rates over the same tenfold increase in revenue.⁴ The exchange rate between conservation and development must therefore be allowed to vary by field office.

The choice situation: one natural feature. The planning record organizes zoning around individual natural features. Figure 2 maps one region of Four Rivers, Idaho, under three configurations the office considered. Panel 2a plots the region’s geothermal zoning, which the plan anchors on four features: a historic trail, a critical groundwater area, elk winter range, and sage-grouse habitat. Panels 2b and 2c map each configuration’s zoning for one feature, the trail and the groundwater area, respectively. The zoning boundaries follow the features themselves. Along the historic trail, the restriction forms a constant-width band around the trail; when zoning protects the critical groundwater area, it stops exactly at the area’s surveyed boundary. In each case, the geometry of the underlying feature dictates the zoning’s geometry. The planning text has the same structure: each configuration is a set of decisions, one for each feature, naming restrictions to protect that feature. The resulting plan is therefore composed of distinct feature-level footprints rather than drawn to a landscape-scale design.

The alternatives: regime–buffer pairs. In the considered configurations, every zoning decision is a regime paired with a buffer, and the two are set separately. When configurations protect a feature differently, they differ only in the regime, the extent, or both. For example, along Figure 2’s historic trail, Four Rivers moved both: the protective band widens from an eighth of a mile to a quarter to a half across the three configurations, and the regime tightens from no surface disturbance to closure in the most restrictive. The set of alternatives at a feature is therefore a cross product: every regime paired with every buffer in the feature’s range.⁵ Reading the zoning for each feature as a regime–buffer pair requires thinking about buffers differently for different kinds of features. Along the trail the buffer is a distance; at a prairie dog town it is the town’s mapped boundary, as at Royal Gorge, Colorado, which prohibits surface disturbance “within active prairie dog colonies



³Townships are 6-by-6 mile squares defined by the Public Land Survey System grid.

⁴This heterogeneity is consistent with the institutional setting: each office writes its own plan for its own planning area while carrying the common mandate to balance development against conservation.

⁵The open regime protects no ground and so pairs with no buffer; Section 3.2 states the set exactly.

as mapped” (Royal Gorge Field Office, Rocky Mountain District 2024, F-16); on a hillside it is a threshold on slope. Diagram 1 draws the first form; Diagram 2, at right, draws the other two. Each form is a single number that grows the protected footprint outward from the feature, so each can be cast as a buffer b on its own scale.

Offices follow several common standards. Figure 3 shows that offices repeatedly choose the same few protection extents. Two distances, a quarter mile and a half mile, account for 85% of protective decisions at raptor nests. Two setbacks, 100 and 500 feet, account for 71% at streams and rivers. Two slope grades, 25% and 30%, account for 64% at steep slopes. These choices are shared across offices. For example, of the ten offices with raptor nests in the sample, eight use quarter-mile buffers and seven use half-mile buffers to protect at least some nests. Two conclusions follow. *First*, the clustering suggests that offices consider more than local benefits and costs when choosing conservation policies. Because development potential varies across features, the same protection extent can forgo very different amounts of revenue. If offices weighed only local net benefits, protection extents would plausibly vary more, both within and across offices. Instead, choices cluster around a few common values. The model therefore allows for a cost of departing from standard practice. *Second*, the multi-modal clustering suggests that offices may follow different standards for the same kind of feature. At raptor nests, for example, 48% of protective decisions use a quarter-mile buffer, while 37% use a half-mile buffer. Because I do not observe which standard an office follows for each feature, I use the median protection chosen by other offices for the same kind of feature as a *proxy*. But a single median cannot capture multiple standards. I therefore allow the standard an office follows to differ from the proxy by an unobserved correction. Without this correction, an office following a quarter-mile standard when the proxy is half a mile would appear to place little weight on standard practice, even though it is following a different standard.

2.3 The Formal Model

This section formally defines the decision problem that Section 2.1 outlines. For each feature f in office i ’s planning area, the office chooses one zoning decision $\mathbf{y} = (r, b)$ ⁶ from the set $\mathcal{V}_{i,f}$ feasible at that feature.⁷ I assume the office’s payoff is additively separable across the features $\mathcal{F}_i$ it administers, so the choice can be represented feature by feature. Specifically, at each feature, the office solves

$$\max_{\mathbf{y} \in \mathcal{V}_{i,f}} U_{i,f}(\mathbf{y}), \quad U_{i,f}(\mathbf{y}) = \overbrace{V_{i,f}(\mathbf{y}; \eta_{i,f})}^{\text{modeled payoff}} + \overbrace{\sigma_{i,k(f)} \varepsilon_{i,f,\mathbf{y}}}^{\text{idiosyncratic shock}}$$

where the subscript $k(f) \in \{\text{soil, water, wildlife}\}$ denotes the feature’s *natural-capital class*. With revenue as the numeraire, I decompose the modeled payoff into four parts:

$$V_{i,f}(\mathbf{y}; \eta_{i,f}) = - \overbrace{M_{i,f}(\mathbf{y})}^{\text{revenue forgone}} + \beta_{i,k(f)} \overbrace{K_{i,f}(\mathbf{y})}^{\text{disturbance avoided}} - \kappa_{i,k(f)} \overbrace{\Delta_{i,f}(b; \eta_{i,f})}^{\text{departure from the standard}} + \overbrace{\gamma_i(r)}^{\text{regime fixed effects}} \quad (1)$$

⁶The policy $\mathbf{y}$ is specific to office i and feature f . I have suppressed the subscripts for clarity.

⁷The pair $\mathbf{y} = (r, b)$ describes the zoning of the fluid leasable resources viable near f : conventional oil and gas, coalbed natural gas, and geothermal.

The two forces of Section 2.1 enter separately: the stakes of the decision through $M_{i,f}$ and $K_{i,f}$, and the pull of the standard the office follows through the departure term $\Delta_{i,f}(b; \eta_{i,f})$.

$M_{i,f}(\mathbf{y})$ is the revenue the policy choice $\mathbf{y}$ forgoes relative to leaving the feature’s surroundings open. It rises as r becomes more restrictive, because a stricter regime lowers expected development and therefore expected revenue, and it rises with b , because the restriction reaches more acres. $K_{i,f}(\mathbf{y})$ is the disturbance to feature f that $\mathbf{y}$ prevents relative to the open baseline. It rises as r becomes more restrictive and rises with b at a *decreasing* rate: disturbance farther from the feature matters less to its natural capital, so each acre added by a wider buffer averts less relevant disturbance. $\Delta_{i,f}$ is the distance between a considered buffer and the standard the office follows: the proxy of Section 2.2, the median buffer the other offices chose for the feature’s family,⁸ plus the correction $\eta_{i,f}$, which the office knows and the econometrician does not.⁹

Revenue forgone enters with a coefficient of one, so the payoff is denominated in dollars. The coefficient $\beta_{i,k(f)}$ is therefore the dollars of forgone revenue office i is willing to give up per unit of relevance-weighted disturbance avoided. One acre of avoided disturbance at full relevance contributes one unit to $K_{i,f}$ and is therefore valued at $\beta_{i,k(f)}$, while an acre farther from the feature is valued at β times its relevance, so its value falls as its relevance declines.¹⁰ β is specific to the feature’s *natural-capital class* $k(f)$, allowing the same office to value the same quantity of avoided disturbance differently depending on what kind of natural capital it protects. The coefficient $\kappa_{i,k(f)}$ is the cost the office pays for departing from the standard, in dollars of forgone revenue. In practice, zoning decisions are developed through a collective decision-making process involving many bureaucrats, stakeholders, and public input. I represent that process in reduced form as a single field office, whose modeled preferences— $\beta_{i,k(f)}$, $\kappa_{i,k(f)}$, and the distribution of $\eta_{i,f}$ —reflect the outcome of those interactions.

Finally, $\gamma_i(r)$ captures systematic preferences for the regulatory regime itself, above and beyond the revenue and disturbance it produces. It distinguishes alternatives that leave land fully open, restrict development without prohibiting it, or close land to development. Stakeholders can make a regime costly or attractive in itself: operators and county governments oppose closure independent of the stakes (Appendix C.4), while conservation groups press the other way, so the same regime can carry a cost at one office and a premium at another. The term $\varepsilon_{i,f,\mathbf{y}}$ is a normalized idiosyncratic choice shock, and $\sigma_{i,k(f)}$ governs its magnitude relative to the modeled payoff. Together, $\sigma_{i,k(f)}\varepsilon_{i,f,\mathbf{y}}$ captures influences on a particular feature-level choice that the modeled payoff does not otherwise represent—the local particulars of the ground, which parties turned up to comment, a powerful local landowner, and other feature- and alternative-specific influences. Section 4 specifies its distribution.

3 Data and Measurement

Section 2 models a field office choosing zoning feature by feature. This section outlines how I derive what the model requires from the planning record. Section 3.1 describes how large language models (LLMs) index, label, and code the written planning record, producing coded values, each linked to the pages of the record that support the coding. Section 3.2 *georeferences* the coded values, placing them on the map, to construct

⁸A feature’s *family* is every feature of its kind, e.g., every perennial stream.

⁹Section 3.3 precisely defines how I construct revenue forgone, disturbance avoided, and the departure from the standard from the planning record.

¹⁰Section 3.3 precisely defines how I model relevance decreasing with distance from the feature.

the choice situation (the feature f), the chosen alternative y^* , and the feasible alternatives $\mathcal{Y}$. Section 3.3 constructs the choice attributes—revenue forgone M , disturbance avoided K , and the departure from the standard Δ —and defines the relevance scale on which K and Δ rest.


3.1 Constructing the Planning-Record Database

This subsection describes how I process the written planning record into a form useful for empirical research. I first *index* each record, constructing a database of its potentially relevant passages; then *label* those passages by what they establish; and finally *code* from the labelled passages what each record states: the regime–buffer pair an office chose, the pairs it considered, and the office’s projections of revenue and disturbance under distinct zoning configurations. These *coded values* do not enter estimation directly: Sections 3.2 and 3.3 describe how I georeference and combine them to construct the estimation sample.

The difficulty in processing the planning record is finding *every* passage that states a regime–buffer pair or a projection to be coded.¹¹ The median office’s record runs to about six thousand pages, the longest to more than twenty thousand, and discussion of any one feature is scattered across the entire record. For raptor nests, say, the buffer an office adopted sits in one appendix, the buffers it considered in a comparison table of a different appendix, and a wider buffer a commenter proposed in a third appendix. Diagram 3 lays out the procedure I use to go from this scattered record to reliably coded values. Its left column shows how the record passes through three readings, each covering fewer pages than the last and each converting the record’s text into a more structured form. Its right column shows how each reading is checked: the first by a literal text search, the second and third by having each agent reread the pages cited in the previous reading rather than trust the earlier agents’ conclusions, and the final coded values by hand.

The first reading covers every page of the record. I divide it among many *agents*: LLMs each given one narrow task. An agent reads about thirty pages at a time searching for discussion of one topic, either a group of related feature families (water and riparian features, say) or one resource (oil and gas), and logs each relevant passage—a stipulation, a configuration considered, the office’s reasoning for its choice—recording the document and page it comes from and a quote

Diagram 3: From Planning Record to Structured Data



Notes: The grey boxes are the successive states of the material, each listing in italics what one of its entries records. Every downward arrow is a reading by language-model agents (filled boxes), and each reading covers fewer pages than the one before. The side column holds the checks; the dashed boxes are the steps where non-language-model text processing enters the procedure: literal text matching and humans. Counts cover feature families and resources together.

¹¹Coding a given value rests on the one or two passages that state it. I nonetheless log every passage that bears on the value, for two reasons. The related passages often provide the context needed to read the passage that states the value. And the passage that correctly states a value is not always the most obvious one, which becomes clear only when all the related passages are read together.

or description of its text.¹² For features alone, every page is therefore read about eight times, once for each group: about 725,000 page-reads over the eleven records’ 90,000 pages; the resource reading brings the total to roughly a million. Each page an agent reads is recorded as read, so coverage of the record can itself be checked. A LLM’s reading can miss a passage, and a miss is silent: a considered regime–buffer pair that is never logged drops out of the feasible set with nothing in the data to show it. I therefore complement the reading with a literal text search that matches strings exactly and involves no LLM: once the record has been read, each topic’s vocabulary in that plan—designation names, stipulation codes, table labels—is collected from the cited passages and searched for across the record. The reading for a topic is not accepted until every page the search flags is either cited or recorded as irrelevant.

A reader familiar with text-processing tools may ask why I have agents read every page rather than begin with semantic retrieval, which represents passages as numerical vectors and uses their proximity to a query to identify likely matches. Beginning with retrieval would be cheaper because agents would need to read only the passages the search selects. But at the outset I do not know most of the vocabulary a plan uses for a topic, and a query written without that vocabulary misses many relevant passages. Reading is itself a search on meaning: rather than rank passages by their similarity to a query, each agent applies a written definition of what counts—for a buffer, a restriction on an activity within a stated distance of a feature—to every page, whatever words the plan uses. Refining the queries over repeated searches does not necessarily close the gap, because the new vocabulary would come only from passages earlier searches happened to return. I instead pay the larger upfront cost of having agents read the full record because omissions matter directly for the data: a passage that is never surfaced may contain a regime–buffer pair an office chose or considered, which would then disappear silently from the feasible set. Full reading also reveals each plan’s own vocabulary—its designation names, stipulation codes, and table labels—across the record rather than only in passages returned by an initial search. I then search that vocabulary exactly, so a passage the agents missed can still be found wherever it uses a term discovered elsewhere in the plan. In ongoing work, I add semantic retrieval as a further audit: the logged passages provide examples of the relevant concepts, and I search the full record for passages that express those concepts in different language but were not logged.

The first reading produces a database of passages, each linked to the feature family or resource it relates to; it holds more than twenty-three thousand passages, about twenty thousand of them on feature families. The reading is deliberately permissive: an agent logs any passage that bears on a feature family or resource, postponing judgement about whether the passage supplies information the estimation needs. An agent that sees thirty pages cannot always judge whether a passage will ultimately matter for coding, because its importance may depend on passages elsewhere in the record. The second reading labels each passage by what it establishes—a leasing restriction on the family, in any configuration the plan adopted or considered; another land-use control on the same ground; or background, a passage that mentions the family without restricting anything. In the third reading, agents turn the labelled passages into model inputs, each agent working only from the passages whose labels fit its task. This keeps each coding task narrow: an agent is given only the passages that bear on the quantity it is asked to recover, not every passage that mentions the feature family or resource. For features, these agents encode each office’s chosen and considered regime–buffer pairs from the passages labelled as restrictions—about half of the passages logged for features. For

¹²Grouping feature families balances cost against mistakes: an agent asked to watch for every family at once misses more, and reading the record once for each single family would multiply the reading several times again.

resources, agents extract the office’s projections of development, revenue, and disturbance from the passages the second reading labelled as containing such projections. The passages left out of a coding task are not discarded: they stay in the database, and when a passage that states a value lacks the context needed to interpret it, the related passages often supply that context.

Several checks stand between a logged passage and a coded value, shown in the right-hand column of Diagram 3. Each stage is run by a different agent that returns to the cited pages instead of trusting a summary from the agent that read the pages in the previous stage, so a passage is read independently several times before it supports a coded value. When a later agent finds that a passage does not match an earlier agent’s description of it, the agent corrects the entry and records the change, so that every correction can be reviewed. When the record itself is ambiguous, the agent stops and puts the question to me: a plan states a restriction but no distance or map for it, a stipulation could belong to either of two feature families, or a plan’s text conflicts with itself. Finally, I personally verify every final coded value against the passages it rests on. Verification establishes two things: that the cited pages support the coded value, and that they come from the documents that record the office’s final decision, not from a draft analysis or a configuration it only considered. Ludwig et al. (2025) show that when an LLM’s output serves as a measurement in a downstream estimate, valid inference requires checking that output against a validation sample. Here I hand-validate the full set of final coded values rather than use a validation sample. Verification does not establish that the reading is complete. A stipulation the full reading overlooked stays overlooked, a page that supports the wrong feature family but reads plausibly passes, and nothing checks that the alternatives recorded are all the alternatives the plan weighed, only that each one recorded exists.

The three-stage procedure is standardized and rerunnable even though an LLM’s reading is not deterministic. The instructions are written down: plain-text files, stored with the replication code, that one orchestrating agent follows step by step, assigning the reading, labelling, and encoding to other agents as it goes.¹³ Two runs will not log exactly the same passages. What is fixed is the standard each must meet: every page the text search flags must be accounted for, and every coded value must cite a page that supports it. The written procedure serves two purposes. First, the same instructions are applied to every office, so all eleven records are read and coded by one procedure. Second, they can be rerun on offices already coded, as an audit: agents read the record afresh, and the fresh reading is reconciled against the stored database. The audit guards against the one failure the text search cannot catch: a passage the first reading missed whose wording appears nowhere else in the record.

Beyond supplying the coded values that enter the estimation sample, the database serves as an index to each record, organized by topic and by label. To review why offices chose the buffers they did for one feature family, for example, an agent can query the passages labelled as background for that family across all eleven offices and select those that argue for or against a given buffer. I can run the identical query and read the same evidence myself. Either way, the question is answered from the database without rereading a record in full.

¹³The instructions call fixed scripts that record each passage and its provenance in the database. Agents reach the documents through a fixed set of tools—search the record, fetch one page, locate a section—so an agent retrieves only the pages a task needs rather than whole documents, which is what makes reading at this scale affordable.

3.2 Constructing Choice Situations and Alternatives

The choice situation. Section 2 defines the natural feature as the choice situation, and speaks of *a* nest or *a* hillside as if each came with its own boundary. In reality, boundaries between features are not well defined. Nests sit close enough together that their buffers overlap, one steep hillside grades into the next, and a stream network runs from one basin into another. Constructing the estimation sample therefore requires a rule for where one feature ends and the next begins. I define a feature at the finest spatial unit the planning record shows offices zoning separately. What that unit is differs by class. For *soil* features—hillslopes, water- and wind-erodible soils, saline soils—a feature is all of one watershed’s ground of that kind, e.g., its hillslopes.¹⁴ Plans zone fragile soils for the waters they drain to, and when they zone a characteristic more finely than the planning area, the split often follows a watershed.¹⁵ For *water* features—streams and rivers, waterbodies, riparian wetlands, springs and wells, drinking-water sources—a feature is likewise everything of one family within one watershed: every perennial stream of the watershed, say, or every spring.¹⁶ For *wildlife* features—raptor nests, prairie dog towns, bald eagle roosts—a feature is a cluster of like sites that share one decision, each within roughly half a mile of another’s protected footprint.¹⁷ Offices frequently zone spatially proximate sites as one unit, and my grouping rule emulates that practice.¹⁸ A reader familiar with land-use plans may object that plans often zone at the level of the feature family, not the instance: every perennial stream in the planning area, not the perennial streams of one watershed. I opt for the finer unit because offices have the discretion to zone at it, and exercise that discretion when it suits them.¹⁹

The sample of choice situations. Field offices select which features to protect from mapped inventories of their planning area. I do not observe each office’s complete inventory, so I define the universe of feature families from the planning record: the union, across all offices, of every family that a considered configuration protected or a public comment raised. No adopted plan therefore determines on its own which feature families enter the sample. The estimation sample covers fifteen families across eleven field offices, grouped by the natural-capital class each protects: four *soil* families (slope, water erodibility, wind erodibility, saline soils), eight *water* families (major rivers, perennial streams, intermittent streams, waterbodies, riparian wetlands, floodplains, drinking-water sources, groundwater), and three *wildlife* families (bald eagle roosts, prairie dog towns, and raptor and other nests).

The planning record contains a broader universe than the estimation sample—plans also anchor restrictions on big-game seasonal ranges, cultural and historic sites, and the historic trail in Figure 2, among others. I restrict attention to the three classes whose underlying inventories can generally be reconstructed and checked. Offices build their soil and water inventories from national surveys—the national soil and elevation

¹⁴Watersheds are the fifth-level hydrologic units (HUC10) of the National Watershed Boundary Dataset.

¹⁵For example, Rock Springs, Wyoming closes the saline soils of the Red Creek watershed while leaving comparably saline soils elsewhere in its planning area open (Appendix C.2).

¹⁶Floodplains and named major rivers instead follow their river systems across watershed lines, so each contiguous floodplain and each named river corridor is its own feature.

¹⁷Concretely, a nest half a mile from the edge of a neighbor’s quarter-mile buffer joins that neighbor’s feature; a third nest half a mile from the second’s buffer joins it too, even if it lies well over a mile from the first.

¹⁸Grand Junction, Colorado, for example, places a stipulation on every occupied prairie dog town in its planning area and a stricter one on the cluster of towns inside a wildlife area it designated (Grand Junction Field Office 2015, B-47).

¹⁹Royal Gorge, Colorado, for example, applies a stricter regime and a stricter buffer to steep slopes inside the South Park Leasing Area than to steep slopes elsewhere in its planning area (Royal Gorge Field Office, Rocky Mountain District 2024, F-25, F-31). Rock Springs, Wyoming and Rio Puerco, New Mexico make the same kind of split, on saline soils and winter elk habitat; Appendix C.2 documents each.

surveys for the soil families, and the national hydrography, wetland, and flood-hazard inventories for the water families—supplemented by state databases where the relevant datum is state-defined.²⁰ I construct the features from the same sources and verify them against the office’s own inventory where it is published and, otherwise, against the acreage the plan reports for each feature family. The wildlife families come from state wildlife agencies’ public inventories—Colorado Parks and Wildlife’s bald eagle nests and roosts, Wyoming Game and Fish’s raptor nests—and, where those are incomplete, from the office’s own databases obtained under the Freedom of Information Act.²¹ The remaining classes are left for future work.

A feature family enters an office’s sample when it is present on ground the office zones, whether or not any considered configuration proposed protecting it. Leaving a present feature open is itself a choice the model prices, and dropping offices that did not consider protecting it would select the sample on the outcome. The obvious alternative rule—admitting a family at an office only when that office’s record protects or mentions it—would in practice produce nearly the same sample: an office’s record discusses almost every feature family present in its planning area, if not in a considered configuration then in a public comment. Gating on presence therefore closely tracks the record while guarding against missed mentions in records that average more than eight thousand pages an office. A family drops out only when there is essentially nothing to decide about: Four Rivers, Idaho has saline soils in its planning area, but almost none on land where the BLM zones minerals, so its saline family is excluded. The wildlife families satisfy the same presence rule, but a nest or a prairie dog town is known to exist only where a state agency or the office has surveyed and recorded it. A wildlife family therefore enters an office’s sample when the available inventory maps at least one site on ground the office zones.²²

The chosen alternative. For each feature in the sample, the chosen regime r is the stipulation the approved plan attaches to the feature, and the chosen buffer b is the extent that stipulation states. Section 3.1 describes how each was read from the record and how the reading was checked. The record does not always state one regime over one extent per feature, and three rules close the gap. First, zoning is written resource by resource, and a plan can assign different resources different restrictions at the same feature (Appendix C.3). I bundle the zoning of all fluid leasable resources into the one decision modeled here; at every feature in the sample, the office assigned those resources the same regime over the same extent. Second, a decision can stack a stricter core inside a weaker ring—closed within a quarter mile of a bald eagle roost and discretionary restrictions out to a mile. I take the stricter core as the chosen alternative and drop the weaker ring.²³ Third, a plan may state the stipulation for a whole feature family—every burrowing owl nest in the planning area—or for one feature, such as the cluster of nests inside a particular management area. I assign each feature the stricter regime–buffer pair that applies to it.

²⁰Wyoming’s stream-quality classes and permitted water wells (State Engineer’s Office), and Idaho’s source-water areas and permitted wells (Departments of Environmental Quality and Water Resources).

²¹Four Rivers, Idaho’s nest and lek layers were obtained this way.

²²Several inventories publish the protected disk rather than the site itself: Colorado’s nests and roosts arrive as quarter-mile disks, and a plan’s own stipulation layers are dissolved and clipped to the land it administers. I recover each site by fitting circles to the arcs of the published boundary, and check each recovery by re-buffering the recovered sites and measuring the share of the published polygon reproduced and the share of the rebuilt polygon that spills outside it. Shrinking the disk back by its radius does not work: overlapping disks have merged, clipped disks have lost the part that would shrink to the centre, and a disk eroded by exactly its own radius is empty.


²³The ring’s ground remains in the feasible set as wider rungs, the ring’s own pair among them, which the office is then modeled as declining. The single-pair representation therefore captures the stricter core exactly but omits the weaker protection outside it.

The alternatives. The considered configurations show two things relevant to constructing the feasible set. The policy instrument is a pair (r, b) , and offices vary its two margins separately (Section 2.2), so the feasible set is the feasible range of regimes paired with the feasible range of buffers. The considered configurations are required to span that set (Section 1, Appendix A), so in principle they bracket the feasible domain of each margin. In practice, an individual plan’s configurations do not always appear to do so, so I take the superset of regimes and buffers considered across plans for each feature family, on the premise that an instrument available to one office administering that family was available to the others. This avoids treating an omission from one plan’s considered configurations as evidence that the omitted choice was infeasible.

For the regulatory regime, any regime is feasible for every office and feature. The planning record contains five: Open, timing limitations, controlled surface use, no surface disturbance, and Closed. Timing limitations and controlled surface use add restrictions relative to Open, limiting when activities may take place over the year and where operations may occur; no surface disturbance bars drilling on the surface entirely but still allows directional drilling from adjacent land; and Closed removes the acreage from leasing altogether. The model collapses these five to three: timing limitations and controlled surface use restrict surface disturbance without prohibiting it and enter as one discretionary regime; no surface disturbance and Closed both prohibit surface disturbance and enter as one closed regime. The regime set is therefore $\mathcal{R} = \{\text{Open, Discretionary, Closed}\}$, with $r \in \mathcal{R}$ at every feature.

For the buffer, the constructed domain is feature-family specific. Open protects no ground, so the range begins at the smallest footprint that protects any ground at all, which depends on the feature’s geometry: zero additional distance for areal geometries such as prairie dog towns, so that the town itself is protected; just above zero for point and line geometries such as nests and streams; and, for threshold families such as slope, the least restrictive threshold, so that only the steepest ground is covered. The maximum of the buffer range is the maximum buffer considered for that feature family across every office’s considered configurations in the estimation sample. Between the minimum and the maximum, I build the feasible buffers as a discrete ladder (Diagram 4): each rung is a dilation of the one below it, and the chosen footprint is itself a rung.²⁴ The ladder preserves the coarseness of the policy choice rather than treating the buffer as continuous. This construction also puts features of the same family with different chosen geometries on one ladder: where an office protected a drawn area rather than a setback, the rungs grow outward from the feature, as they would toward a setback, until they fill the boundary the office drew, so a ring around each site

Diagram 4: The Buffer Ladder



Notes: Each shaded band is one rung of the buffer ladder for a nest, drawn on the grid. The ladder starts at the cell holding the nest (min), grows one cell per step in every direction including diagonals, reaches the chosen footprint, itself a rung, and continues to the ceiling (max).

²⁴Concretely, the ladder starts at the minimum footprint (the open alternative aside), dilates step by step until it reaches the chosen footprint, which is emitted as a rung, and then continues dilating to the maximum buffer when the chosen footprint is not already the maximum. The chosen footprint enters at the distance the plan states; only the rungs between it and the ends of the ladder are grid dilations rather than stated distances. Each step is a fixed number of 8-adjacency dilations on a grid of roughly 30×30 m cells. Threshold families step through successive attribute thresholds instead of dilating. For these families the chosen threshold is the cutoff the plan states, as Royal Gorge, Colorado does for slopes over 30 percent (Royal Gorge Field Office, Rocky Mountain District 2024, F-25, F-31); where the plan instead names a mapped class, “fragile soils of high concern” for water erosion, the cutoff on the underlying soil variable is read from the plan’s own class definition or calibrated so that the class reproduces the acreage the plan cites for it.

and the boundary drawn around all of them sit on the same ladder.²⁵

Collecting the two margins, the feasible set is $\mathcal{Y}_{i,f} = \{\text{Open}\} \cup (\mathcal{R} \setminus \{\text{Open}\}) \times \mathcal{B}(f)$, where $\mathcal{R}$ is common to every office and feature, and $\mathcal{B}(f)$ is the ladder just described: its range is fixed by f 's feature family, and its rungs are footprints of f itself. The set is finite, and it is what the likelihood in Section 4 sums over. The threat to this construction is over-inclusion, not omission: an alternative feasible somewhere may not have been feasible everywhere—a statutory floor under a listed species, a ceiling this office could not defensibly draw. What limits the concern is that the set's bounds come from configurations some office actually considered (Appendix D.1). An alternative infeasible at a given office would therefore have to be so for a reason specific to that office. The most common such reason—that a wider footprint reaches only ground the office does not zone—is handled by the attribute construction: attributes are measured only on the office's own land, so a rung that adds no additional zoned ground is identical to the rung below and is dropped as a duplicate.

3.3 Constructing the Choice Attributes

I construct the three attributes—revenue forgone, disturbance avoided, and the departure from the standard—from the office's own analysis: directly where the plan quantifies the stakes, and by modeling the qualitative judgments it records where it does not. Regulation requires every plan to compare its considered configurations in a benefit–cost framework, projecting the development each would permit and the natural capital it would protect (Appendix A). I use these projections to measure the tradeoffs offices perceived among their options. For the choice model, these perceived stakes are the relevant objects: the aim is not the most accurate forecast of revenue and disturbance, but the expectations, accurate or not, that informed the zoning decision.

To construct revenue forgone and disturbance avoided at the feature level, I replicate the office's accounting at a finer spatial scale. Offices map development potential across the planning area and intersect each zoning configuration with that map to project wells, production, revenue, surface disturbance, and protected acreage relative to an open baseline (Figure A1 reproduces Rock Springs, Wyoming's ledger). Where the office performs this calculation for an entire zoning configuration, I perform it for a single feature's policy footprint, once for each feasible choice.²⁶ The one input I model rather than read from the plan is how much an acre's disturbance matters to the feature it lies near. Plans routinely treat an acre of disturbance as less consequential the farther it lies from the feature, but express that relationship in words rather than numbers. I encode that judgment on the *relevance scale* Section 2.2 measures choices on, so disturbance avoided depends not only on the acreage protected but on where that acreage lies.

Costs: revenue forgone. Let $M^s(\mathbf{y})$ be the expected revenue *forgone* from resource s , one of the fluid resources viable near f , under the restrictions $\mathbf{y}$ imposes, relative to the open ($r = \text{Open}$) baseline.²⁷ The office's projection for resource s starts from a map that partitions the planning area into development-

²⁵Casper, Wyoming, for example, protects most of its bald eagle roosts as the mapped roost stands, but around the roosts in Jackson Canyon it drew a mapped area, extending the boundary to the topographic features that screen the roosts (Casper Field Office, Wyoming State Office 2007, 2-41).

²⁶One concern is that these tabulations could themselves be constructed to justify a choice already made. Appendix D.1 discusses how the public planning process constrains this possibility.

²⁷The office and feature indices are suppressed throughout this subsection: every attribute is built for one feature f at one office i , and the $M(\mathbf{y})$ and $K(\mathbf{y})$ defined here are Section 2's $M_{i,f}(\mathbf{y})$ and $K_{i,f}(\mathbf{y})$.


potential categories $c \in \mathcal{C}^s$ —low, medium, and high in Diagram 5.²⁸ Each category carries a *development rate*: the number of wells the office expects per acre if that land is left open. Multiplying this rate by the office’s expected revenue per well gives category c a per-acre revenue density $m_{i,c}^s$.²⁹ A more restrictive regime permits only a share $a_{r,c}^s$ of the wells the baseline projects in category c —stipulations raise operating costs, a closed regime permits no wells on the footprint—so the regime removes the share $1 - a_{r,c}^s$ of the baseline’s wells, and with them the revenue they would earn. Revenue forgone is therefore:

$$M^s(\mathbf{y}) = \sum_{c \in \mathcal{C}^s} \overbrace{(1 - a_{r,c}^s)}^{\text{Regime scaling in category } c} \overbrace{m_{i,c}^s}^{\text{Footprint's acreage in category } c} \underbrace{|B_f(b) \cap c|}_{\text{Policy footprint}}, \quad (2)$$

where the *policy footprint* $B_f(b)$ is the area over which regime r applies. So, M^s is the footprint’s acreage in each category, valued at that category’s density and scaled by the development the regime removes. Diagram 5 shows the intersection the sum runs over: a single footprint, and the several potential categories its acreage falls in.³⁰ M^s rises with r because $a_{r,c}^s$ falls with r , and rises weakly with b because the footprint’s acreage increases.³¹

Section 2’s revenue forgone attribute sums this per-resource quantity M^s over the fluid resources viable near f —each with its own potential map, densities and shares, and the same $\mathbf{y}$ applied to all of them: $M(\mathbf{y}) = \sum_{s \in \mathcal{S}} M^s(\mathbf{y})$ where $\mathcal{S}$ is the set of resources viable near f .³²

Diagram 5



The relevance scale. The relevance scale assigns each distance l from feature f a value $w_f(l)$: full weight at the feature itself, declining as l grows, and negligible beyond an *effective ecological range* h :

$$\underbrace{w_f(0) = 1}_{\text{full weight at the feature}}, \quad \underbrace{w'_f(\cdot) \leq 0}_{\text{declining}}, \quad \underbrace{w_f(h) \approx 0}_{\text{negligible by the range}}.$$

²⁸See Figure A1a for an example from the planning record. Figure A1a plots Rock Springs’s oil and gas development-potential map.


²⁹This revenue is the gross value of the production the office projects. It accrues to operators and reaches federal and state treasuries only through royalties, rents, and bonus bids—it is not revenue the office itself collects, and the office faces no budget constraint. What the estimates below price is therefore the value an office places on the development and fiscal interest, not a marginal utility of income.

³⁰The intersections $|B_f(b) \cap c|$ have no closed form—footprints trace the landscape and routinely straddle several categories—and are computed by intersecting the mapped footprint with the potential map, category by category. Appendix F details how the densities—revenue and disturbance alike—are recovered from Reasonably Foreseeable Development (RFD) reports.

³¹Specifically, the shares are ordered $1 = a_{\text{Open},c}^s \geq a_{\text{Discretionary},c}^s \geq a_{\text{Closed},c}^s = 0$.

³²Summing across resources treats each resource’s expected revenue as additive: developing one resource on an acre is assumed not to preclude developing another there, and because each resource earns its own revenue stream, the sum then double counts nothing. Where two resources are in fact mutually exclusive on the same ground, the sum overstates the revenue at stake; the additive form assumes such ground is rare, and I leave richer interactions among co-located resources to future work.

Any function with these three properties would serve; the one I use is $w_f(l) = 0.01^{(l/h)^p}$, one at the feature and a hundredth of that at the range h , whatever the curvature p .³³ The range h is *where* relevance becomes negligible, and the curvature p *how* it declines to that point. Diagram 6, at right, draws the curve at $p = 2$ and shades the band the curvatures $p \in [1.5, 2.5]$ sweep out. Evaluated at the buffer's edge, $w_f(b)$ is the relevance of the last acre the buffer protects, which the diagram's dotted lines read off the curve; Section 2.2 measures buffer choices on this scale.³⁴ Because relevance runs from 1 to 0 at every feature, whatever the geometry of its policy footprints or its range h , the relevance of the last protected acre puts every feature's feasible buffers on one scale, so departures from the standard compare across all features.



The planning record supplies the underlying judgment that relevance declines with distance (Appendix C.5), but not its exact functional form. The range h is fixed per feature family from ecological evidence outside the planning record, and held constant across offices rather than fit to the buffer widths planners chose, which would confound it with the very preferences the model sets out to recover.³⁵ The curvature p has no such grounding and is not estimated; Section 5 reports the estimates under two curvatures, $p = 1.5$ and $p = 2$.

Benefits: disturbance avoided. Let $K^s(\mathbf{y})$ be the expected surface disturbance *relevant to* f that the restrictions avoid, relative to the open baseline. Multiplying the development rate by the office's expected surface disturbance per well gives category c a per-acre disturbance density $d_{i,c}^s$. Unlike a dollar of revenue, worth the same wherever earned, an acre of disturbance avoided next to an owl's nest protects the nest far more than one a mile away. Each 30×30 m grid cell j therefore enters weighted by its relevance $w_f(l_j)$, with l_j its distance from the feature:³⁶

$$K^s(\mathbf{y}) = \sum_{c \in \mathcal{C}^s} (1 - a_{r,c}^s) d_{i,c}^s \overbrace{\sum_{j \in B_f(b) \cap c} w_f(l_j) |j|}^{\text{Relevance-weighted acreage in category } c}. \quad (3)$$

K^s rises with r as $a_{r,c}^s$ falls, and rises with b in proportion to the relevance the added acres carry, which is negligible past the range h . Section 2's disturbance avoided attribute sums K^s over the fluid resources

³³For prairie dog towns and bald eagle roosts, whose sites are areas rather than points, an acre's relevance is instead the convolution of the site's footprint with $w_f(\cdot)$, normalized so that its maximum is one. For a point site this reduces exactly to the $w_f(l)$ of the text. For an area, it weights each acre by how much habitat lies within reach of a disturbance on that acre. An acre inside a compact town, or among clustered towns, therefore carries more relevance than one in an isolated strip. The premise that an acre's relevance depends on the extent and configuration of the habitat around it has support in the planning record: Appendix C.6 collects offices writing town size into their rules and defending them in those terms.

³⁴On the grid, the buffer's edge is read as the 99th percentile of distance over the footprint, which is robust to the few cells a dilation pins against the window; for the two convolved families $w_f(b)$ is the mean relevance over the footprint's boundary cells.

³⁵Appendix E lists each family's range and derives it: a few hundred metres for water features, a few kilometres for wildlife, and, for the soil families, whose features carry no geometry to measure distance from, a range in attribute space rather than metres.

³⁶For the two areal wildlife families the cell's relevance is instead the normalized convolution of footnote 33. For the threshold families, whose scale is in attribute space, l_j is how far the cell's grade, erodibility, or salinity falls below the level at which relevance is full, and h the shortfall at which it becomes negligible (Appendix E).

viable near f , $K(\mathbf{y}) = \sum_{s \in \mathcal{S}} K^s(\mathbf{y})$.³⁷

The departure from the standard. The last attribute measures how far an alternative’s buffer lies from the standard the office follows for the feature: the proxy, the median buffer chosen elsewhere in the agency, plus the feature’s correction. Let $w_f(b)$ be, as above, the relevance of the last acre a feasible buffer b protects, and $\bar{w}_{-i,f}$ the *proxy* for f ’s feature family: the median $w_f(b)$ *chosen* for features of that family across the *other* field offices, leaving office i out so that an office’s own decisions never move the yardstick they are measured against. The departure of equation 1 is the squared gap between an alternative’s $w_f(b)$ and the standard the office follows,

$$\Delta_{i,f}(b; \eta_{i,f}) = \left[\overbrace{w_f(b)}^{\text{last acre's relevance}} - \underbrace{\left(\overbrace{\bar{w}_{-i,f}}^{\text{proxy}} + \overbrace{\eta_{i,f}}^{\text{correction}} \right)}_{\text{standard the office follows}} \right]^2, \quad (4)$$

where $\eta_{i,f}$ is the correction. At $\eta_{i,f} = 0$ the office follows the proxy itself; a negative $\eta_{i,f}$ puts the standard past the proxy, a wider more protective buffer, and a positive $\eta_{i,f}$ puts it short of the proxy, a narrower less protective buffer.³⁸ $\eta_{i,f}$ is known to the office but not to the econometrician; I specify assumptions about its distribution in Section 4. The other two components, $w_f(b)$ and $\bar{w}_{-i,f}$, come from the relevance scale and the record. The departure is defined only for alternatives that protect the feature: leaving it open sets no buffer, so the open alternative carries no departure, $\Delta_{i,f} = 0$. Measuring the departure as a squared distance from a reference point follows spatial theories of choice (Davis et al. 1970); centering that reference point on a social norm while allowing offices to follow different standards is motivated by Bernheim (1994).

4 Identification and Estimation

Section 2 sets out the payoff a field office maximizes, and Section 3 describes how I construct the data to estimate the model. This section first explains what variation in the data identifies the office’s valuation of disturbance avoided relative to the revenue numeraire, and what variation distinguishes that valuation from the cost of departing from the standard. It then specifies the distribution of the correction that locates the standard an office follows relative to the proxy, and the distribution of the idiosyncratic shock. Together these distributions turn the payoff model into a probabilistic choice model, yielding the likelihood used for estimation.

4.1 Identification

³⁷Summing disturbance across resources assumes their disturbed acres do not coincide: no shared pads, roads, or other infrastructure. Where two resources are developed on the same ground, the sum over-counts the disturbance at stake; as for revenue, the additive form assumes such ground is rare.

³⁸As explained in Section 2.2 and at Diagram 6, relevance decreases as protection extends outward, so a negative $\eta_{i,f}$ lowers the standard on the relevance scale and makes it more protective.


What identifies the Valuation of Disturbance Avoided relative to the Revenue Numeraire.

Revenue forgone and disturbance avoided are built from the same ground. As Diagram 5 and equation 2 show, revenue forgone is an acreage-weighted sum of the per-acre revenue densities $m_{i,c}^s$ over the policy footprint. Without the relevance weight $w_f(l_j)$, equation 3 would be the same acreage-weighted sum of the disturbance densities $d_{i,c}^s$. The two densities ($m_{i,c}^s$ and $d_{i,c}^s$) are proportional: each is the category's development rate multiplied by a per-well amount, revenue for one and surface disturbance for the other. $M_{i,f}$ and an unweighted $K_{i,f}$ would therefore be nearly proportional across a feature's alternatives. If the proportionality were exact, $K_{i,f} = \lambda M_{i,f}$, the payoff from the two attributes would collapse to $(\beta_{i,k(f)} \lambda - 1) M_{i,f}$. Every extension of the buffer would then change the payoff by $(\beta_{i,k(f)} \lambda - 1) \Delta M_{i,f}$, i.e., the two attributes move as one, leaving no variation to identify $\beta_{i,k(f)}$ separately from the scale of the idiosyncratic shock, $\sigma_{i,k(f)}$. Appendix Table A2 shows how close that proportionality is: with all acres weighted equally, the median within-feature correlation between revenue forgone and disturbance avoided is 1.000, and the 5th percentile is 0.983. Differences across regimes add little independent variation in $M_{i,f}$ and $K_{i,f}$: among alternatives that share a regime and differ only in the buffer, the 5th percentile is 0.986, barely above the 0.983 across all alternatives.

The relevance weight reduces this collinearity, allowing $\beta_{i,k(f)}$ to be identified. It enters disturbance avoided and never revenue forgone, so as protection extends outward, $M_{i,f}$ and $K_{i,f}$ change at different rates. Diagram 7, above at right, builds both attributes on Diagram 5's nest and potential map: the two panels plot expected revenue forgone and disturbance avoided as the policy footprint grows from nothing to beyond the nest's effective ecological range h . Revenue forgone climbs at an increasing rate, as the footprint's area grows and its outer rings reach higher potential ground, while the marginal disturbance avoided by added acres falls as the added acres' relevance decreases with b . As the buffer approaches and passes the range h , those added acres carry almost no relevance, so disturbance avoided all but stops increasing with b . That falling rate is what makes the office's stopping point informative. Setting aside the cost of departure and the idiosyncratic shock, an office extends the buffer as long as the added disturbance avoided is worth the added revenue forgone, $\beta_{i,k(f)} \Delta K_{i,f} > \Delta M_{i,f}$, and stops when it is not. Because $\Delta K_{i,f} / \Delta M_{i,f}$ falls toward zero as the buffer approaches h , an office that stops early reveals a low $\beta_{i,k(f)}$, and one that protects out to the range reveals a high one. With relevance-weighted acres, the median correlation in Table A2 falls to 0.868 and the 10th percentile to 0.669; among alternatives that share a regime, to 0.847 and 0.591.

This argument parallels Hendel (1999), who notes that in a standard discrete-choice model estimated on a single cross-section, brand-specific dummies would absorb the variation in price, because prices vary at the brand level, leaving the price coefficient unidentified. His multiple-unit demand model avoids that problem because firms choose both the quality and the number of computers: expenditure is linear in the units purchased, while the payoff is a non-linear function of their quantity and quality. Here a single choice, the extent of the buffer, sets both how many acres an office protects and their relevance to the feature, the counterparts of Hendel's quantity and quality. Revenue forgone is roughly linear in the acres protected, while disturbance avoided is non-linear in them because each added acre carries less relevance than the last. Variation in where offices stop along that schedule carries information about $\beta_{i,k(f)}$.

Diagram 7



What separates the Cost of Departure from the Valuation of Disturbance Avoided. Both coefficients act on the buffer choice: a positive $\kappa_{i,k(f)}$ pulls the buffer toward the standard $\bar{w}_{-i,f} + \eta_{i,f}$, while a positive $\beta_{i,k(f)}$ pushes it outward because disturbance avoided $K_{i,f}$ is valuable. What separates them is that the attributes they multiply depend on different things: the departure $\Delta_{i,f}$ on where the buffer sits on the feature’s relevance curve, disturbance avoided $K_{i,f}$ on the ground beneath the policy footprint. The departure depends on the buffer through $w_f(b)$, the relevance of the last acre protected: among features with the same standard, two buffers at the same point on their respective relevance curves incur the same departure, regardless of the development potential beneath their footprints. Disturbance avoided depends on the ground the buffer lands on: two features with the same standard and the same $w_f(b)$, and therefore the same departure, can score K far apart because the development their regimes prevent differs. $K_{i,f}$ and $\Delta_{i,f}$ therefore do not move in lockstep across features, providing the independent variation that identifies the two coefficients separately.

4.2 Distributional Assumptions

Two parts of an office’s payoff are not observed: the correction that locates its standard relative to the proxy, and the idiosyncratic differences between choices’ payoffs. This subsection assumes a distribution for each. I assume a discrete support for the correction $\eta_{i,f}$, common to all offices and feature families, with office-specific mixing shares, and a type I extreme value distribution for the idiosyncratic shock. Each choice probability is then a finite mixture of logits.

The correction and its distribution. The correction $\eta_{i,f}$ locates the standard the office follows relative to the proxy $\bar{w}_{-i,f}$. The office observes η , and I do not: it knows which standard it is following, whether or not that is the proxy. As discussed in Section 2.2, η lets an office follow a different standard at different features, and different offices follow different standards. A standard is by definition a protection extent shared across many features, so a discrete distribution for η is the natural choice. Assuming a discrete support amounts to a finite set of standards for each kind of feature. Specifically, I assume that $\eta_{i,f}$ takes one of a fixed set of values, common to every field office and feature family: $\eta_{i,f} \in \{\eta_1, \dots, \eta_Q\}$. Because $\eta_{i,f}$ is measured on the relevance scale, which runs from 1 to 0 at every feature whatever its family’s range or geometry, the same support values can apply to all features, even though a given η_q implies a different buffer at each family. A shared support therefore does not mean the same standards across families: each standard is the family’s proxy plus the correction. Sharing that support across offices and feature families lets the full sample locate the $\{\eta_q\}$. Pooling disciplines the fit by preventing the model from assigning a separate standard to nearly every choice. The tradeoff is the assumption that standards can be more or less protective than their proxies by similar amounts across feature families, measured on the common relevance scale. This is plausible if offices draw on guidance from state agencies or earlier plans that prescribe similarly stricter or more permissive protections across different kinds of features.

Alongside the Q corrections, each office carries one more: an open type that leaves the feature unprotected and therefore takes no position on where protection should stop. This separates an office’s propensity not to protect at all from the standard it follows when it does protect. Without the open type, the Q corrections would have to account for both margins, so open choices would also influence the parameters intended to describe the office’s standard conditional on protecting. The open type faces neither the cost of departure

nor the development–conservation tradeoff—it simply walks away.³⁹ The *mixing shares are office-specific*, and run over all $Q + 1$ types:

$$\Pr(\eta_{i,f} = \eta_q \mid i) = s_{iq}, \quad q = 1, \dots, Q, \quad s_{i0} + \sum_{q=1}^Q s_{iq} = 1,$$

where s_{i0} is the open-type share.

Choice probabilities and likelihood. The idiosyncratic shock is $\sigma_{i,k(f)}\varepsilon_{i,f,\mathbf{y}}$, where $\sigma_{i,k(f)}$ sets its scale and $\varepsilon_{i,f,\mathbf{y}}$ is a type I extreme value draw, independent across the alternatives in a feature’s choice set and across features. Given type q , the choice probability takes the multinomial logit form whose payoff is equation 1 evaluated at $\eta_{i,f} = \eta_q$:

$$P_{i,f}(\mathbf{y} \mid \eta_q) = \frac{\exp\{V_{i,f}(\mathbf{y}; \eta_q)/\sigma_{i,k(f)}\}}{\sum_{\mathbf{y}' \in \mathcal{Y}_{i,f}} \exp\{V_{i,f}(\mathbf{y}'; \eta_q)/\sigma_{i,k(f)}\}}.$$

The unconditional probability mixes those Q logits with a point mass on the open type:

$$P_{i,f}(\mathbf{y}) = s_{i0} \mathbf{1}[r = \text{open}] + \sum_{q=1}^Q s_{iq} P_{i,f}(\mathbf{y} \mid \eta_q),$$

where the open type leaves the feature open with probability one. The Q types also leave features open when the stakes and the departure make that the best alternative; s_{i0} therefore captures an additional propensity to leave features unprotected beyond what that payoff predicts. The valuations $\beta_{i,k}$, cost-of-departure parameters $\kappa_{i,k}$, scale parameters $\sigma_{i,k}$, closure fixed effects $\gamma_i(\text{closed})$ ⁴⁰, common support $\{\eta_q\}$, and office-specific shares s_{iq} maximize the log likelihood of the observed policies,

$$\sum_i \sum_{f \in \mathcal{F}_i} \log P_{i,f}(\mathbf{y}_{i,f}^*).$$

Because the mixture has finite support, the likelihood is evaluated exactly rather than by simulation. The number of types Q is set by adding support points until the next one carries a negligible share; Section 5 reports the fitted support.

4.3 Estimation Restrictions and Units

Estimation departs from equation 1 in two ways: κ and σ are pooled to a coarser grouping, and the payoff is estimated in log-odds units rather than dollars. It also holds the closure effect fixed across an office’s groups.

κ and σ are estimated on a coarser grouping. Above, both κ and σ vary with the natural-capital class $k(f) \in \{\text{soil, water, wildlife}\}$. In estimation, I instead allow them to vary over two groups: soil and

³⁹With the open type introduced, I drop the fully-open policy fixed effect of equation 1, so office-specific residual variation in the propensity to leave a feature unprotected is carried by the open type rather than by a fully-open policy fixed effect.

⁴⁰The remaining regime effect once the fully-open effect is dropped, measured relative to the discretionary regime.

water together, and wildlife separately; write this coarser grouping $g(f)$. When soil and water are estimated separately, some office–class combinations contain too little variation to pin down the scale and cost-of-departure parameters reliably, and estimates can run toward 0 or ∞ . Pooling soil and water stabilizes the estimates and is a natural restriction because the two classes are measured largely on the same watershed units (Section 3.2).

The model is estimated in log-odds units. With the pooling just described, the idiosyncratic shock is $\sigma_{i,g(f)}\varepsilon_{i,f,\mathbf{y}}$. The choice probability depends on the payoff only through $V_{i,f}/\sigma_{i,g(f)}$, so I estimate in those units, with the revenue coefficient $1/\sigma_{i,g(f)}$ a free parameter. Conditional on type q , the estimated utility is therefore

$$-\frac{1}{\sigma_{i,g(f)}} M_{i,f} + \theta_{i,k(f)} K_{i,f} - \phi_{i,g(f)} \Delta_{i,f}(b; \eta_q) + \gamma_i(\text{closed}) \mathbf{1}[r = \text{closed}] + \varepsilon_{i,f,\mathbf{y}},$$

where $\theta_{i,k(f)} = \beta_{i,k(f)}/\sigma_{i,g(f)}$ and $\phi_{i,g(f)} = \kappa_{i,g(f)}/\sigma_{i,g(f)}$. I convert the estimated coefficients into willingness-to-pay space, $\beta_{i,k(f)} = \sigma_{i,g(f)}\theta_{i,k(f)}$ and $\kappa_{i,g(f)} = \sigma_{i,g(f)}\phi_{i,g(f)}$, with standard errors computed by the delta method.

One closure effect per office. I estimate the closure fixed effect with one $\gamma_i(\text{closed})$ shared across an office’s groups. Allowing that effect to vary by group leaves some office–group combinations too weakly identified, with estimates running toward positive or negative infinity. The common effect therefore enters as a fixed shift in the log-odds of closing, measured relative to the discretionary regime. Because $\sigma_{i,g(f)}$ still varies across groups, its implied dollar equivalent, $\sigma_{i,g(f)}\gamma_i(\text{closed})$, differs across an office’s groups. Section 5 accordingly reports β and κ in revenue units and the closure fixed effect in log-odds units.

5 Results

This section reports what the fitted model shows. Every office is willing to give up revenue to conserve, though how much varies widely. Offices also tie their choices to conventions that do not depend on the benefits or costs of conserving a feature. They pay to stay near the standard: at the median office, the buffer the model predicts forgoes a higher net benefit—conservation value net of revenue forgone—at an adjacent buffer in seven of ten decisions. And each office is habitually inclined toward closing a feature outright or against it, whatever the net benefits.

Willingness to pay. In the raw data, Figure 1 shows offices conserve less as the revenue at stake rises, at visibly different rates. The model puts a price on this pattern: each office’s willingness to pay (WTP) for conservation, the revenue it gives up to avoid disturbing a feature, by non-market resource. Figure 4 plots the WTPs, in dollars of forgone revenue per acre of avoided disturbance at the feature, with one estimate for each office and non-market resource. Table 1 reports the same estimates and their standard errors, alongside the closure fixed effect and the cost of departure. Twenty-eight of the thirty-two willingness-to-pay estimates are positive and statistically significant at the ten percent level. At the median office, WTP is \$58,534 per acre for soil, \$66,425 for water, and about \$37,600 for wildlife. The spread across offices is wide. Specifically, WTP for soil ranges from \$22,983 per acre at Uncompahgre, CO to \$406,979 at Royal Gorge, CO at ten of the eleven offices; at the eleventh, Little Snake, CO, the estimate is indistinguishable from zero (−\$3,990,

standard error \$6,907). For water, WTP ranges from \$12,038 at Rio Puerco, NM to \$1,605,124 at Royal Gorge, CO, and it is positive and significant at all eleven offices. For wildlife, WTP ranges from \$26,966 at Casper, WY to \$897,196 at Royal Gorge, CO at seven of the ten offices that carry wildlife families; at Four Rivers, ID (\$793, \$4,824), Grand Junction, CO (\$3,630, \$26,853) and Lander, WY (\$65,302, \$66,252) the estimate is smaller than its own standard error.

The pattern of positive and statistically significant WTP estimates survives two specification checks. *First*, Figure A4 plots estimates from the same mixed logit fit with a relevance curvature of $p = 1.5$; the main specification uses $p = 2$. Across the two curvatures, all thirty-two WTP estimates keep their sign, the same twenty-eight are statistically significant, and magnitudes are similar: every $p = 1.5$ estimate lies within a factor of 1.9 of its $p = 2$ counterpart. *Second*, Table A1 and Figure A3 report estimates from a conditional logit: the same model with the standard departure term and the open type removed, so that an office is modeled as a pure benefit-cost optimizer. The two models agree on the sign of thirty-one of the thirty-two WTP estimates; the exception is Little Snake soil, one of the four mixed-logit estimates indistinguishable from zero. Magnitudes differ more than signs, and mostly in one direction: the mixed logit’s estimates are larger. For soil, the two models’ estimates are within 30% of each other at eight of the eleven offices. For water, the mixed logit produces the larger estimate at nine of the eleven offices, and for wildlife at seven of the ten, with differences as large as a factor of seven (Buffalo wildlife, \$321,292 versus \$43,653). Omitting the departure cost explains part of the difference. In the mixed logit, departing from a standard is allowed to be costly; the conditional logit sees no standard, let alone a cost of departing from it. It reads a buffer at a standard as if the benefits and costs put it there: a standard short of the buffer the office’s true value would pick on benefits and costs alone is read as a low value on conservation, and one past it as a high value. The mixed logit instead credits that buffer to the standard. And when a buffer does move past the standard, the office pays the departure cost as well as the revenue, so the mixed logit estimates a higher WTP. This mechanism explains why the wildlife estimates differ most: there, standards sit short of where the benefits and costs of conservation would put the buffer.

The willingness-to-pay estimates vary with local politics in the direction one would expect. Areas with more Trump voters are plausibly more supportive of development on public land, and offices in those areas are willing to pay less to protect non-market resources. Figure 5 plots each office’s willingness to pay against the share of the 2016 presidential vote cast for Donald Trump in its planning area. The relationship is not causal: the vote share only stands in for an area’s political leaning. Rio Puerco, NM is the exception. Its planning area gave Trump 43% of the vote, the lowest share in the sample, but its willingness to pay for water, \$12,038 per acre, is also the lowest. Weighting each precinct by its acres of BLM land, rather than counting every vote in the planning area once, sharpens both patterns (Figure A5). The other ten offices’ shares all move toward Trump, by 8 to 22 percentage points, and the negative relationship among them steepens; Rio Puerco’s share moves 6 points the other way, and it stands further apart. Rio Puerco, NM is plausibly an exception because New Mexico is one: its voters lean Democratic, but the state depends heavily on revenue from the federal onshore mineral estate (New Mexico Legislative Finance Committee 2023).⁴¹ Rio Puerco’s low Trump share therefore overstates local support for restricting development.

⁴¹New Mexico’s share of federal royalties is paid into the state’s general fund and came to \$2.3 billion in fiscal year 2022 (New Mexico Legislative Finance Committee 2023). When the Biden administration paused new oil and gas leasing on federal land in 2021, New Mexico’s two Democratic senators asked the White House to compensate the state for the revenue it stood to lose; Volcovici & Renshaw (2021) describe New Mexico as “by far the biggest beneficiary of the federal drilling program.”

The standard and the cost of departing from it. The data support including a cost for departing from standards. Adding it to the conditional logit discussed above raises the log-likelihood by about 8,000 points for 93 parameters. Ninety-three parameters that explained nothing would raise it by about 45, so this gain is more than a hundred times what chance alone would produce. The rest of this section outlines where the fitted model locates the standards and the implied costs an office faces for departing from them.

In the planning record, offices defend their buffers by citing standards: a wildlife agency’s recommendation, a state office’s guidance, or a previous decision (Appendix D). I do not consistently observe which standard an office anchors its choice on, so I proxy for the standard with the median choice of the other ten offices, $\bar{w}_{-i,f}$. Figure 3 shows the proxy works reasonably well: about a fifth of protective decisions sit on it. But decisions also cluster systematically short of the proxy and beyond it, so I allow a correction η to help locate an office’s standard. The fitted model estimates six corrections. Figure A6’s top section shows they straddle the proxy, from $\eta_1 = -0.691$ to $\eta_6 = +0.430$, with one at -0.002 sitting on the proxy itself. A positive correction puts the office’s standard short of the proxy, a narrower buffer; a negative correction puts it past the proxy. The same section prints, under each correction, the buffer it implies for raptor nests: the proxy protects a nest to 760 m; the two positive corrections put the standard at 430 m ($\eta = +0.29$) and 230 m ($+0.43$); the correction at -0.16 puts it at 940 m; and the two most negative corrections (-0.52 and -0.69) put it beyond 1.5 miles (2,414 m), the end of the nest’s ecologically relevant range. The correction at $+0.29$ therefore moves a nest’s standard from about a half mile to about a quarter mile, a buffer eight of the ten offices apply to at least some raptor nests. The six corrections imply six standards for nests, but not six standards in all: a standard is the family’s proxy $\bar{w}$ plus a correction η , so the same correction implies a different buffer for a raptor nest than for a stream.

If the corrections locate standards, which exist outside any one office and are available to all offices, two predictions follow, and both hold in the fitted model. *First*, there are only so many standards to cite, so a few corrections should be enough. The fit finds six suffice. The number of corrections is not set in advance: as Section 4 describes, I add corrections one at a time until the next one carries a negligible share. The search stops at six, for eleven offices and 135 office-by-feature family pairs, because a seventh would carry only 2% of the mixing shares, all of it at one office, Grand Junction.⁴² *Second*, offices should rely on the same corrections. If a correction instead captures an office’s own idiosyncratic taste for conservation, each office would have a correction of its own, and nothing would make those tastes fall on a few shared points instead of spreading over a continuum. Consistent with this prediction, offices do rely on the same corrections. Figure A6’s rows plot each office’s mixing shares over the six corrections: the estimated probability that the office anchors a decision to the standard that correction implies.⁴³ The model estimates that six of the eleven offices anchor a decision to the correction $+0.29$ with a probability of at least 50%. Three offices anchor to the correction -0.16 with a probability of at least 25%. Five anchor to the proxy itself with probabilities ranging from 2% to over 60%. Three offices anchor to the correction $+0.43$ and two to the corrections -0.52 and -0.69 . Averaged across the eleven offices (Figure A6’s bottom row), the mixing share is 12% on the proxy itself, 60% on the two corrections short of it, 27% on the three past it, and 2% on the open type.⁴⁴

One might object that the corrections are few and discrete by construction: I assume a discrete distribution,

⁴²Going from five corrections to six raises the log-likelihood by 89 points, and going from six to seven by 39 points, each for 12 added parameters: one η and eleven office shares.

⁴³The shares are estimated through a softmax, which can approach zero or one but never reach them. The shares the figure leaves blank are below 0.1%, and all but one of them are 10^{-11} or smaller; a circle marked 100% is one to the same precision.

⁴⁴The open-type share s_{i0} is below one percent at eight of the eleven offices; Colorado River Valley carries the most, at 8%.

a shared support forces offices onto common points, and the stopping rule caps how many corrections there are. The model fit contests this objection. If decisions were spread smoothly across the range, as they would be if there were no standards to cluster around, a handful of corrections would approximate the spread by spacing themselves evenly across it, and each added correction would move the others to cover the range more finely. The fit shows neither. First, the corrections are not evenly spaced: three sit past the proxy, at -0.69 , -0.52 and -0.16 , one sits on it, and two sit short of it, at $+0.29$ and $+0.43$, with gaps between neighbors from 0.14 to 0.36. Second, the corrections do not move as new ones are added. Going from five corrections to six, the new correction sits at -0.69 , beyond the others; the corrections from the five-correction fit move by at most 0.01. Similarly, going from six to seven, five of the six stay within 0.003 and the largest move is 0.04. In summary, the fit confirms what the planning record suggests: offices anchor their choices on standards, and this confirmation is not by construction.

The cost of departure, κ , is the revenue an office gives up to stay close to those standards. The fit finds offices do pay to stay close, in two senses: the costs are not zero, and they are large enough that predicted choices often forgo an alternative with a higher net benefit to stay closer to the standard. *First*, the cost cannot be negative in estimation but it can be zero; only one of the twenty-one estimates sits at zero, and eighteen are distinguishable from zero at the ten percent level. Table 1’s final four columns report the cost of departure (κ), for soil and water together and for wildlife. The first and third of them give κ itself: like β , it is in dollars of forgone revenue, here per unit of the squared gap Δ between the chosen buffer and the standard the office anchors on, $\bar{w}_{-i,f} + \eta_{i,f}$. The second and fourth price the departure each office actually makes: the annual dollar-equivalent cost on an average protective decision: for each such decision, the squared gap between the chosen buffer and each type’s standard, weighted by that type’s posterior probability given the choice, averaged over the office’s protective decisions and multiplied by κ .⁴⁵ Priced at the departure each office actually makes, the soil-and-water penalty runs from \$3,750 at Rio Puerco and \$31,266 at Four Rivers to \$1,070,296 at Little Snake and \$2,059,150 at Grand Junction. On the wildlife side it runs from \$6,174 at Uncompahgre and \$6,450 at Grand Junction to \$86,143 at Casper and \$2,120,810 at Buffalo. Lander’s wildlife penalty is larger still but imprecise: its standard error exceeds the estimate. At the other extreme, Four Rivers bears no cost of departure for wildlife: the estimate sits at $\kappa = 0$, the edge of the parameter space, and with the other parameters held at their estimates the likelihood falls for every positive value, with a one-sided 95% bound of \$36 per unit of squared gap.⁴⁶

Second, predicted choices often forgo an alternative with a higher net benefit. Figure A7 plots the share of each field office’s decisions in which a buffer adjacent to the *predicted* buffer has a higher net benefit—its conservation value net of the revenue given up. The figure includes the decisions for which the fitted model predicts protection: with the latent types weighted by the office’s mixing shares, it assigns the protective alternatives together a higher probability than leaving the feature open. The predicted buffer is the protective alternative with the highest predicted probability. Its adjacent buffers are the alternatives under the same regime one step narrower and one step wider. Adjacent buffers differ only in net benefit and in the cost of departure, so without a cost of departure no adjacent buffer would have a higher net benefit than the

⁴⁵The posterior probability of a type given a choice is the office’s mixing share on that type times the probability the type assigns to the chosen alternative, normalized across types; on a protective choice, the only kind that enters this average, the open type receives no weight.

⁴⁶The bound is the value of κ at which the log-likelihood, with every other parameter held at its estimate, has fallen 1.353 points below its value at $\kappa = 0$, the one-sided 95% likelihood-ratio cutoff ($\chi^2_1(0.90)/2$); the bound is one-sided because κ cannot be negative.

predicted one, and every share would be zero. Instead, at the median office an adjacent buffer has a higher net benefit in 69% of decisions: 71% for soil, 66% for water and 75% for wildlife. Across offices the share runs from 44% at Royal Gorge to 94% at Buffalo. In each of these decisions the cost of departure holds the predicted buffer in place: moving to the adjacent buffer would raise the cost of departure by at least the gain in net benefit. Table 2 shows that the gain in net benefit from an adjacent buffer is frequently minimal, but sometimes quite large. At the median office the gain exceeds \$100 a year in 49% of decisions, \$1,000 in 32% and \$10,000 in 12%. The large gains are concentrated in a few offices: the gain exceeds \$10,000 a year in 77% of decisions at Buffalo and 64% at Little Snake, against 3% at Four Rivers and Rio Puerco.

The regime margin. The cost of departing from standards captures an office’s habit of picking particular buffers. The closure fixed effect $\gamma_i(\text{closed})$ captures a second habit: how much more or less inclined an office is, whatever the net benefits, to close a feature outright rather than restrict development with a discretionary stipulation. It is negative at the four Wyoming offices and at Rio Puerco—Rock Springs most sharply, at -7.6 log-odds—positive at the remaining six, and significant at every office (Table 1). As with willingness to pay, the habit follows state politics. The five Colorado offices, in a state that leans toward conservation, are inclined to close; the four Wyoming offices, in a state that leans toward development, are inclined not to. To show how much this habit matters for predicted choices, Figure A8 plots, for each office, the share of protective decisions that close the feature rather than stipulate it: as realized, as the fitted model predicts, and as the model predicts with the office’s closure fixed effect set to zero. Without the fixed effect, every office but Little Snake would close between 46% and 56% of the features it protects. The fixed effect is what moves Rock Springs from 52% predicted closure without it to 0.3% with it—three closures among 1,035 protected features—and Colorado River Valley from 49% to 84%.

In conclusion, offices respond to the benefits and costs of conservation, and they tie their decisions to standards, treating a departure from standards as a cost. BLM headquarters could therefore change what offices protect in two ways: by recommending a specific level of protection (setting a standard), or by changing the value offices place on avoided disturbance (setting a value for conservation). Section 6 compares the implications of these two policy instruments.

6 Policy Implications

In this section, I bring the estimates together to examine how economic tradeoffs and established conservation standards jointly shape allocation decisions. I first consider a counterfactual in which every BLM office responds to the Fish and Wildlife Service’s highly protective one-mile buffer recommendation for birds’ nests, holding all other model primitives fixed. This exercise quantifies how adopting a more protective recommendation at BLM headquarters could change conservation outcomes and costs, measured by additional birds protected and revenue forgone. I then ask whether a more traditional price-like instrument, which sets a common minimum value for avoiding disturbance around nests, could achieve the same conservation gains at lower cost. Finally, I compare the cost-effectiveness of both policy instruments to a benchmark: the outside option cost of saving a bird through retrofitting power poles.

6.1 Revising the standard

The first counterfactual moves the nesting standard to the Fish and Wildlife Service’s (FWS) one-mile recommendation. The FWS’s primary responsibility is wildlife conservation, unlike the BLM, which must balance development and conservation under its multiple-use mandate. Its recommendation therefore provides a more conservation-oriented standard for BLM offices. This counterfactual asks what would happen if BLM headquarters encouraged offices to justify their choices against the FWS recommendation. This contrasts with existing BLM guidance to use “the least restrictive [policy] that effectively accomplishes the resource objectives” (Bureau of Land Management 2013).

I consider two implementations of this counterfactual, *heterogeneous adoption* and *uniform adoption*, reflecting different assumptions about how offices would respond to new guidance from BLM headquarters. Under *heterogeneous adoption*, I assume $\eta_{i,f}$ reflects persistent differences due to local practices or ecological conditions, so offices may continue to follow different conservation standards despite receiving the same recommendation. I therefore set the nesting proxy to one mile, i.e., $\bar{w}_{-i,f} = w_f(1 \text{ mile})$, while retaining the office- and nest-specific correction $\eta_{i,f}$. Under *uniform adoption*, I assume $\eta_{i,f}$ reflects differences in the guidance available to offices at the time of decision-making, so the recommendation eliminates these differences and all offices respond to the same standard. I therefore set the entire nesting standard to one mile, i.e., $\bar{w}_{-i,f} + \eta_{i,f} = w_f(1 \text{ mile})$ for all offices and nests. Because I cannot empirically determine which assumption is more accurate, I report counterfactual results under both.

I measure conservation gains in expected baby birds (fledglings) protected, rather than acres restricted. Following FWS guidance, my accounting assumes a nest produces two baby birds in the absence of nearby drilling. I use the plans’ forecasts of drilling near each nest to calculate the expected number of fledglings under each alternative, assuming independent well placements. Figure A9 plots the changes in baby birds protected and revenue forgone under the two implementations of the revised standard, relative to the baseline. Under heterogeneous adoption, approximately 1,300 additional fledglings are protected at a cost of \$337 million in revenue forgone, or \$254,000 per additional fledgling. Under uniform adoption, approximately 240 more fledglings are protected for less than \$10 million in additional revenue forgone, reducing the average cost to \$220,000 per fledgling. The difference between the two implementations illustrates how the cost of departing from standard practice can leave relatively inexpensive birds unprotected while keeping relatively expensive birds protected. Under heterogeneous adoption, the office- and nest-specific correction $\eta_{i,f}$ shifts the standard offices anchor on relative to the common one-mile recommendation. In Buffalo, WY, this correction makes the standard less protective than the recommendation, while in Casper and Lander, WY, it generally makes the standard more protective (see Figure A6). Under uniform adoption, these corrections are eliminated, making Buffalo’s standard more protective and those of Casper and Lander less protective than under heterogeneous adoption. Consequently, Buffalo protects more birds, while Casper and Lander protect fewer. The additional birds protected in Buffalo are relatively inexpensive to protect. Its less protective standard, rather than high revenue costs, had discouraged their protection. Conversely, the birds no longer protected in Casper and Lander were relatively expensive to protect. These offices had protected them despite their high revenue costs because reducing protection would have incurred large departure costs. Uniform adoption therefore redirects conservation toward birds that are less costly to protect, increasing conservation while reducing its average cost.

This counterfactual is a back-of-the-envelope exercise intended to illustrate how revising conservation

standards could change allocation decisions. I use a one-mile buffer as a deliberately protective interpretation of FWS guidance for nests in the estimation sample. The FWS recommends one mile for nests of unidentified raptors, ferruginous hawks, and bald eagles, but one-half mile for golden eagles, falcons, and goshawks and one-quarter mile for Swainson’s and red-tailed hawks (U.S. Fish and Wildlife Service 2022). Many nests in the planning record lack species identification, and ferruginous hawks predominate among identified nests. Thus, the one-mile buffer matches FWS guidance for much of the sample but exceeds it for some identified species. The results therefore represent the conservation gains and revenue costs of applying this protective one-mile standard uniformly, rather than following FWS guidance species by species.

6.2 Setting the value of conservation

The second policy counterfactual changes the value offices place on avoiding wildlife disturbance, holding all else fixed, including the nesting standard. This counterfactual considers a more traditional price-like policy instrument. Rather than recommending a buffer distance, BLM headquarters would prescribe a dollar value for avoided wildlife disturbance, much as the federal government has historically prescribed a social cost of carbon for agencies to use in their analyses. In the model, this intervention changes $\beta_{i,\text{wildlife}}$, the annual revenue office i is willing to forgo to avoid one relevance-weighted acre of wildlife disturbance.

I consider two implementations of this counterfactual, a *minimum value* and a *common value*, reflecting different assumptions about how headquarters would treat offices that already value wildlife more than it recommends. Under a *minimum value*, headquarters sets a floor $\underline{\beta}$: offices with valuations below this floor must raise them to $\underline{\beta}$, while offices already above it retain their valuations, i.e., $\beta_{i,\text{wildlife}}^{\text{cf}} = \max\{\beta_{i,\text{wildlife}}, \underline{\beta}\}$. Under a *common value*, headquarters requires all offices to use the same valuation, raising some offices’ valuations and lowering others.

I calibrate both implementations to protect 1,327 additional baby birds as the revised standard under heterogeneous adoption. This allows me to compare the costs of achieving the same conservation gains under the two types of policy instruments. Calibrating the valuations is straightforward because the number of additional baby birds protected increases monotonically with the prescribed valuation. Under the *minimum value* implementation, setting the lower bound on offices’ valuations at \$212,000 per relevance-weighted acre protects 1,327 additional baby birds, the conservation gains under the revised standard. Under the *common value* implementation, setting all offices’ valuations to \$240,000 per relevance-weighted acre achieves the same conservation gains. As expected, the calibrated common value is higher than the minimum value because it also reduces valuations in offices initially above the prescribed value, causing them to protect fewer birds. Therefore, to achieve the same aggregate conservation gains requires greater protection in offices with lower initial valuations.

Figure 6 plots the changes in fledglings protected and revenue forgone under the two price-like instruments alongside the revised standard. Both price-like instruments achieve the same aggregate conservation gains at approximately the same revenue cost as the revised standard, but differ substantially in where birds are protected. Under the standard revision, Buffalo accounts for 680 of the 1,327 additional fledglings protected. Under the minimum value counterfactual, Buffalo protects no additional birds, while Rock Springs, Casper, and Lander, WY, protect far more than under the standard revision. This difference arises because Buffalo already values avoiding wildlife disturbance at more than \$321,000 per relevance-weighted acre, substantially higher than every office except Royal Gorge, CO (see Figure 4). A price-like instrument would therefore

need to prescribe a value above \$321,000 to induce additional protection in Buffalo. At that value, however, offices like Rock Springs, Casper, and Lander would protect far more fledglings than needed to achieve the aggregate conservation gains of the revised standard. Under the common value counterfactual, the contrast is even starker. Buffalo protects even fewer birds, while Rock Springs, Casper, and Lander protect slightly more than under the minimum value counterfactual. These differences in where birds are protected explain why the revised standard is as cost-effective as the price-like instruments. As discussed above, Buffalo has many birds that could be protected at relatively little revenue cost but remain unprotected because of its less protective standard. Revising that standard induces protection of these relatively inexpensive birds. In contrast, a common valuation cannot induce protection of these birds without also inducing substantially more protection in other offices. Thus, the cost of departing from standard practice can make a revised standard as cost-effective as a price-like instrument, despite the latter’s theoretical ability to direct conservation toward offices where it is less costly. This finding holds under both assumptions about how offices respond to headquarters’ recommendation. When calibrated to match the 1,566 additional fledglings protected under uniform adoption, both price-like instruments again concentrate additional protection in Casper, Rock Springs, and Lander rather than Buffalo, making them no more cost-effective than the revised standard (see Figure A10).

6.3 Alternative mitigation as a benchmark

The preceding comparisons establish that revising a conservation standard can be at least as cost-effective as prescribing a social value. However, neither instrument is necessarily the least-cost way to protect birds. To assess this, I compare their costs with an alternative form of mitigation. The FWS estimates that offsetting the expected death of one bird (raptor in this region) costs approximately \$120,000, equivalent to retrofitting sixteen power poles to prevent electrocutions (U.S. Fish and Wildlife Service 2024). Applying the FWS’s resource-equivalency accounting to disturbed nests implies mitigation costs of \$91,000 to \$130,000 per baby bird (fledgling) (U.S. Fish and Wildlife Service 2018*b*, U.S. Fish and Wildlife Service 2018*a*). Measured against this benchmark, each policy instrument forgoes approximately twice as much revenue per additional baby bird as the estimated cost of mitigation through power-pole retrofits. For example, protecting 1,327 additional baby birds annually under the revised standard costs \$337 million in forgone revenue, compared with approximately \$159 million for equivalent mitigation at \$120,000 per bird.

I also use this benchmark to examine how conservation would change if BLM headquarters prescribed \$120,000 as the value of protecting a baby bird, implementing the common value counterfactual from the previous subsection at the FWS benchmark. To implement this counterfactual, I translate the value per fledgling into the model’s units of dollars per relevance-weighted acre. Following the FWS buffer guidance, my accounting treats drilling a well at a nest as costing a brood of two fledglings. Thus, a value of \$120,000 per fledgling implies that offices should be willing to forgo \$240,000 to prevent a well from being drilled at a nest. The expected surface disturbance per well ranges from 2 to 16 acres across field offices. I therefore divide \$240,000 by the expected disturbance per well for each office to obtain the corresponding valuation per relevance-weighted acre.

The left panel of Figure A11 applies the reverse conversion to express offices’ estimated willingness to pay in dollars per fledgling, allowing for a more intuitive comparison with the FWS benchmark. Six of the ten offices have estimated willingness to pay above the benchmark, two have estimates near it, and only two

have estimates meaningfully below it. This suggests that setting a common value of \$120,000 per fledgling would reduce conservation. The right panel confirms this result, plotting the change in fledglings protected when all offices are assigned this common value, relative to the baseline (gray). Buffalo protects more than 500 fewer fledglings, while the largest increase in protection at any office is just 12 fledglings. Thus, relative to the FWS mitigation benchmark, existing allocation decisions produce more conservation than a common social value of \$120,000 would induce. Most strikingly, Buffalo accounts for the largest conservation gains under a revised standard but also the largest losses when offices are assigned the FWS mitigation value.

These comparisons are back-of-the-envelope exercises. They apply the FWS’s mitigation cost for a golden eagle to nests of every raptor species in the sample, using the golden eagle as a high-end benchmark. They are intended to provide intuition for how offices’ estimated valuations, average costs of protection under the two policy instruments, and mitigation costs compare, not to establish whether offices’ allocation decisions are socially efficient. A key concern, particularly when using the mitigation benchmark as a common social value, is whether it understates the social value of protecting a bird. However, other available valuation estimates provide limited grounds for prescribing a substantially higher value.⁴⁷ That said, the benchmark-based conclusions are sensitive to assumptions about the effectiveness of alternative mitigation. The FWS’s \$120,000 benchmark assumes power-pole retrofits remain effective for 30 years. If retrofits instead last 10 years, the estimated mitigation cost rises to approximately \$250,000 per fledgling, comparable to the \$254,000 per fledgling under the revised standard. Using this higher benchmark as a common social value also leaves conservation nearly unchanged, rather than substantially reducing it. Thus, while the revised standard remains at least as cost-effective as the price-like instruments under both assumptions about adoption, conclusions about their costs relative to alternative mitigation depend on the assumed effectiveness of that mitigation.

7 Conclusion

This paper studies how the Bureau of Land Management (BLM) allocates federal land between development and conservation, a question that has received little attention from economists (Libecap 2026). I estimate how much development revenue local offices are willing to forgo for conservation. I find that their choices depend on both the local benefits and costs of conservation and adherence to standard practice.

The paper builds on work that recovers public decision makers’ objectives from their choices (Bodéré et al. 2026, Kashner 2026, Slattery 2025, Parkhomenko 2023). My approach differs in how I construct the alternatives available to each office and the information it had when deciding. Federal regulations require BLM offices to document these alternatives and their forecasts. I use large language models to process roughly 90,000 pages of planning records from 11 local offices. The resulting dataset contains 12,509 zoning decisions. For each decision, I recover the alternatives considered and use the office’s own forecasts to calculate the expected revenue forgone and disturbance prevented under each. This allows me to recover the tradeoffs offices faced when deciding, rather than infer them from outcomes years later.

At the median office, willingness to pay for conservation is about \$60,000 in annual revenue per acre of

⁴⁷O’Brien & Hagerty (2025) provide two alternative per-eagle valuation benchmarks: \$127,000 based on costs imposed through criminal enforcement, including fines and compensatory mitigation, and \$22,100 obtained by allocating stated willingness to pay for the existence of bald eagles across the population. Neither directly estimates the marginal social value of an additional fledgling.

disturbance prevented. This value varies across offices and across resources within each office. It is also correlated with local politics. Offices whose planning areas gave Donald Trump a larger share of the 2016 presidential vote are generally willing to give up less revenue for conservation. One interpretation of this correlation is that Trump vote share proxies for local preferences for development over conservation.

Benefits and costs alone do not explain offices' choices. On the extent margin, offices treat departures from standard practice as costly. At the median office, the buffer with the highest predicted choice probability yields lower net benefits, conservation value net of revenue forgone, than an adjacent buffer in 70% of decisions. Offices also differ in how willing they are to close land to development, even after accounting for the benefits and costs of conservation. All four Wyoming offices tend to avoid closure, while all five Colorado offices tend to favor it. This follows the same broad political pattern as willingness to pay for conservation, with Republican-leaning Wyoming less inclined to prohibit development than Democratic-leaning Colorado. These tendencies make a large difference. Without them, the model predicts that all but one office would choose closure in roughly half their allocation decisions. With these tendencies, closure rates range from 0.3% in Rock Springs, WY, to 84% in Colorado River Valley, CO.

Offices' adherence to standard practice, in addition to their response to the benefits and costs of conservation, has two implications for policy. First, BLM headquarters can steer local decisions by recommending conservation standards. Second, adherence to standard practice can make price instruments less cost-effective because some offices stick to less protective standards even when protecting more would be worth the cost. I demonstrate these implications by comparing i) a recommended one-mile no-drilling buffer around raptor nests with ii) price instruments that specify a dollar value for avoiding wildlife disturbance, either as a floor under offices' own valuations or as a common value for all. Counterfactuals show that the recommended standard can be as cost-effective as either price instrument in achieving the same conservation gains. Revising the standard can encourage inexpensive conservation that a price instrument reaches only by inducing more costly protection elsewhere. That a recommended standard can be as cost-effective as a price instrument is striking, but it does not establish that either policy is the best way to protect birds. The Fish and Wildlife Service can protect a bird for about \$120k, roughly half the cost per bird under either land-use policy.

More generally, my findings suggest that how standards develop, spread, and influence future decisions deserves further study, as their influence may be far greater than any one decision. For example, BLM headquarters already uses programmatic plans to steer decisions across local offices. These plans focus on a particular resource or activity, such as protecting sage grouse habitat or facilitating wind development. Debate over programmatic plans tends to focus on the allocations they change when adopted. My findings suggest that their influence extends beyond these initial changes, as offices revising their own plans treat programmatic plan's allocations as standards they must justify departing from.

References

- Angelsen, A., Brockhaus, M., Sunderlin, W. D. & Verchot, L., eds (2012), *Analysing REDD+: Challenges and Choices*, CIFOR, Bogor, Indonesia.
- Aspelund, K. M. & Russo, A. (forthcoming), ‘Additionality and asymmetric information in environmental markets: Evidence from conservation auctions’, *American Economic Review*.
- Barbier, E. B. & Burgess, J. C. (2001), ‘The economics of tropical deforestation’, *Journal of Economic Surveys* **15**(3), 413–433.
URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-6419.00144>
- Becker, G. S. (1983), ‘A theory of competition among pressure groups for political influence*’, *The Quarterly Journal of Economics* **98**(3), 371–400.
URL: <https://doi.org/10.2307/1886017>
- Bernheim, B. D. (1994), ‘A theory of conformity’, *Journal of Political Economy* **102**(5), 841–877.
- Bodéré, P., Brancaccio, G. & Kang, K. (2026), Inefficiencies in environmental infrastructure: Evidence from drinking water in California, Technical report. Working paper.
- Buffalo Field Office, High Plains District (2015a), Buffalo field office approved resource management plan, Technical Report BLM/WY/PL-15/022+1610, U.S. Department of the Interior, Bureau of Land Management, Buffalo, Wyoming.
- Buffalo Field Office, High Plains District (2015b), Proposed resource management plan and final environmental impact statement for the buffalo field office planning area, Technical Report BLM/WY/PL-15/014+1610, U.S. Department of the Interior, Bureau of Land Management, Buffalo, Wyoming.
- Bureau of Land Management (1985), Energy and mineral resource assessment (blm manual section 3031), Technical report, U.S. Department of the Interior. BLM Manual Rel. 3-115, supersedes Rel. 3-52 (June 19, 1985).
URL: https://www.blm.gov/sites/blm.gov/files/uploads/mediacenter_blmmanual3031.pdf
- Bureau of Land Management (2008), H-1790-1 national environmental policy act handbook, Technical report, U.S. Department of the Interior, Bureau of Land Management, Washington, D.C.
URL: https://www.blm.gov/sites/blm.gov/files/uploads/Media_Library_BLM_Policy_Handbook_h1790-1.pdf
- Bureau of Land Management (2013), H-1624-1 planning for fluid mineral resources, Technical report, U.S. Department of the Interior, Bureau of Land Management, Washington, D.C. Release 1-1749.
URL: <https://www.blm.gov/sites/blm.gov/files/H-1624-1%20rel%201-1791.pdf>
- Bureau of Land Management (2023), Changes to the four rivers field office proposed resource management plan, Plan amendment, U.S. Department of the Interior, Bureau of Land Management. Separates oil and gas from geothermal management actions and focuses oil and gas leasing in high or moderate potential areas by closing areas with low or no oil and gas potential.
URL: https://eplanning.blm.gov/public_projects/?doc=/1250/200020763/20077691/250083873/FRFO%20Changes%20
- Bureau of Land Management, Colorado State Office (2020), ‘Response to the governor’s consistency review of the uncomphgre field office proposed resource management plan and final environmental impact statement’, Letter to Governor Jared Polis, January 8, 2020. U.S. Department of the Interior, Bureau of Land Management, Lakewood, Colorado.

- Casper Field Office, Wyoming State Office (2007), Record of decision and approved casper resource management plan, Technical report, U.S. Department of the Interior, Bureau of Land Management, Casper, Wyoming. Updated with amendments and maintenance actions.
URL: <https://eplanning.blm.gov/eplanning-ui/project/68150/510>
- Cochrane, M. A., Moran, C. J., Wimberly, M. C., Baer, A. D., Finney, M. A., Beckendorf, K. L., Eidenshink, J. & Zhu, Z. (2012), ‘Estimation of wildfire size and risk changes due to fuels treatments’, *International Journal of Wildland Fire* **21**(4), 357–367.
URL: <https://doi.org/10.1071/WF11079>
- Collins, C. P. & Reynolds, T. D. (2005), Ferruginous Hawk (*Buteo regalis*): A Technical Conservation Assessment, Technical report, USDA Forest Service, Rocky Mountain Region, Species Conservation Project, Rigby, Idaho. Prepared by TREC, Inc. 64 pp.
URL: https://web.archive.org/web/2018/https://www.fs.usda.gov/Internet/FSE_DOCUMENTS/stelprdb5182004.pdf
- Colorado River Valley Field Office (2014), Colorado river valley field office proposed resource management plan and final environmental impact statement, Technical report, U.S. Department of the Interior, Bureau of Land Management, Silt, Colorado.
- Cropper, M. L., Evans, W. N., Berardi, S. J., Ducla-Soares, M. M. & Portney, P. R. (1992), ‘The determinants of pesticide regulation: A statistical analysis of epa decision making’, *Journal of Political Economy* **100**(1), 175–197.
URL: <https://doi.org/10.1086/261812>
- Davies, K. W., Vavra, M., Schultz, B. & Rimbey, N. (2014), ‘Implications of longer term rest from grazing in the sagebrush steppe’, *Journal of Rangeland Applications* **1**, 14–34.
URL: <https://research.fs.usda.gov/treesearch/download/47536.pdf>
- Davis, O. A., Hinich, M. J. & Ordeshook, P. C. (1970), ‘An expository development of a mathematical model of the electoral process’, *American Political Science Review* **64**(2), 426–448.
- Dixit, A. (2002), ‘Incentives and organizations in the public sector: An interpretative review’, *The Journal of Human Resources* **37**(4), 696–727.
URL: <http://www.jstor.org/stable/3069614>
- Doherty, K. E., Maestas, J., Remington, T., Naugle, D., Boyd, C., Wiechman, L., Bedrosian, G., Cahill, M., Coates, P., Crist, M., Holdrege, M., Kumar, A., Mozelewski, T., O’Connor, R., Olimpi, E., Olsen, A., Prochazka, B., Reinhardt, J., Smith, J., Sparklin, W., Theobald, D. & Wollstein, K. (2024), ‘State of the sagebrush: Implementing the sagebrush conservation design to save a biome’, *Rangeland Ecology & Management* **97**.
- Edwards, E. C., O’Grady, T. & Jenkins, D. (2019), Sooner or safer? bureaucracy in oil and gas production, CEnREP Working Paper 19-017, Center for Environmental and Resource Economic Policy, Raleigh, NC.
- Egan, L. M. & Watts, M. J. (1998), ‘Some costs of incomplete property rights with regard to federal grazing permits’, *Land Economics* **74**(2), 171–185.
URL: <http://www.jstor.org/stable/3147049>
- Ferraro, P. J. & Pattanayak, S. K. (2006), ‘Money for nothing? a call for empirical evaluation of biodiversity conservation investments’, *PLoS Biology* **4**(4), e105.
- Fischel, W. A. (2001), *The Homevoter Hypothesis: How Home Values Influence Local Government Taxation, School Finance, and Land-Use Policies*, Harvard University Press, Cambridge, MA.

- Four Rivers Field Office, Boise District (2019), Four rivers field office draft resource management plan and draft environmental impact statement, Technical report, U.S. Department of the Interior, Bureau of Land Management, 3948 Development Avenue, Boise, Idaho, 83705.
URL: <https://eplanning.blm.gov/eplanning-ui/project/1250/510>
- Frakes, M. (2013), ‘The impact of medical liability standards on regional variations in physician behavior: Evidence from the adoption of national-standard rules’, *American Economic Review* **103**(1), 257–276.
- Frank, E. G., Auffhammer, M., McLaughlin, D., Spiller, E. & Sunding, D. L. (2025), The cost of species protection: The land market impacts of the endangered species act, Working Paper 33352, National Bureau of Economic Research.
URL: <http://www.nber.org/papers/w33352>
- Grand Junction Field Office (2015), Grand junction proposed resource management plan and final environmental impact statement, Technical Report DOI-BLM-CO-N030-2008-0001-RMP-EIS, U.S. Department of the Interior, Bureau of Land Management, Grand Junction, Colorado.
- Hanson, L. A. & Vincent, C. H. (2020), Federal land ownership: Overview and data, Technical Report R42346, Congressional Research Service.
URL: <https://crsreports.congress.gov/product/pdf/R/R42346>
- Hendel, I. (1999), ‘Estimating multiple-discrete choice models: An application to computerization returns’, *The Review of Economic Studies* **66**(2), 423–446.
- Jakus, P. M. & Akhundjanov, S. B. (2019), ‘The antiquities act, national monuments, and the regional economy’, *Journal of Environmental Economics and Management* **95**, 102–117.
URL: <https://www.sciencedirect.com/science/article/pii/S0095069618305849>
- Johnson, C. W. & Buffler, S. (2008), Riparian buffer design guidelines for water quality and wildlife habitat functions on agricultural landscapes in the Intermountain West, Technical Report General Technical Report RMRS-GTR-203, U.S. Department of Agriculture, Forest Service, Rocky Mountain Research Station, Fort Collins, CO.
- Jones, A. & Carter, J. G. (2016), ‘Implications of longer term rest from grazing in the sagebrush steppe: an alternative perspective’, *Journal of Rangeland Applications* **3**, 1–8.
URL: <https://thejra.nkn.uidaho.edu/index.php/jra/article/view/18>
- Joskow, P. L. & Noll, R. G. (1981), Regulation in theory and practice: An overview, in G. Fromm, ed., ‘Studies in Public Regulation’, The MIT Press, pp. 1–78.
URL: <http://www.nber.org/chapters/c11429>
- Kashner, Z. (2026), Building with externalities: Local governments and wind farms, Technical report, Stanford Graduate School of Business. Working paper.
- Lander Field Office, Wind River/Bighorn Basin District (2013), Lander proposed resource management plan and final environmental impact statement, Technical report, U.S. Department of the Interior, Bureau of Land Management, 1335 Main Street, Lander, Wyoming, 82520.
URL: <https://eplanning.blm.gov/eplanning-ui/project/18602/510>
- Leaver, C. (2009), ‘Bureaucratic minimal squawk behavior: Theory and evidence from regulatory agencies’, *American Economic Review* **99**(3), 572–607.
- Lewis, E. (2019), ‘Patchwork policies, spillovers, and the search for oil and gas’, *American Economic Journal: Economic Policy* **11**(1), 380–405.
URL: <https://www.aeaweb.org/articles?id=10.1257/pol.20160373>


- Libecap, G. D. (2026), The federal lands: An economic property rights perspective, Working Paper 35745, National Bureau of Economic Research.
URL: <http://www.nber.org/papers/w35745>
- Little Snake Field Office (2010), Little snake proposed resource management plan and final environmental impact statement, Technical Report FES-10-28, U.S. Department of the Interior, Bureau of Land Management, Craig, Colorado.
- Little Snake Field Office (2011), Record of decision and approved resource management plan for public lands administered by the bureau of land management little snake field office, Technical report, U.S. Department of the Interior, Bureau of Land Management, Craig, Colorado.
- Ludwig, J., Mullainathan, S. & Rambachan, A. (2025), Large language models: An applied econometric framework, Working Paper 33344, National Bureau of Economic Research.
- Magat, W. A., Krupnick, A. J. & Harrington, W. (1986), *Rules in the Making: A Statistical Analysis of Regulatory Agency Behavior*, 1 edn, RFF Press. Republished in 2011.
URL: <https://doi.org/10.4324/9781315064529>
- Martin, W. E., Shields, D. J., Tolwinski, B. & Kent, B. (1996), ‘An application of social choice theory to u.s.d.a. forest service decision making’, *Journal of Policy Modeling* **18**(6), 603–621.
URL: <https://www.sciencedirect.com/science/article/pii/S0161893895001328>
- McCubbins, M. D., Noll, R. G. & Weingast, B. R. (1987), ‘Administrative procedures as instruments of political control’, *Journal of Law, Economics, & Organization* **3**(2), 243–277.
- McCubbins, M. D. & Schwartz, T. (1984), ‘Congressional oversight overlooked: Police patrols versus fire alarms’, *American Journal of Political Science* **28**(1), 165–179.
- McFadden, D. (1974), Conditional logit analysis of qualitative choice behavior, in P. Zarembka, ed., ‘Frontiers in Econometrics’, Academic Press, New York, pp. 105–142.
- McFadden, D. (1975), ‘The revealed preferences of a government bureaucracy: Theory’, *The Bell Journal of Economics* **6**(2), 401–416.
URL: <http://www.jstor.org/stable/3003236>
- McFadden, D. (1976), ‘The revealed preferences of a government bureaucracy: Empirical evidence’, *The Bell Journal of Economics* **7**(1), 55–72.
URL: <http://www.jstor.org/stable/3003190>
- Muhn, J. & Stuart, H. R. (1988), *Opportunity and Challenge: The Story of the BLM*, U.S. Department of the Interior, Bureau of Land Management, Washington, D.C.
URL: <https://www.blm.gov/sites/default/files/documents/files/opportunity%20and%20challenge.pdf>
- New Mexico Legislative Finance Committee (2023), New mexico royalties and state finances, Technical report, New Mexico Legislative Finance Committee. June 28, 2023.
- Oates, W. E. (1972), *Fiscal Federalism*, Harcourt Brace Jovanovich, New York.
- Oates, W. E. & Portney, P. R. (2003), Chapter 8 - the political economy of environmental policy, in K.-G. Mäler & J. R. Vincent, eds, ‘Environmental Degradation and Institutional Responses’, Vol. 1 of *Handbook of Environmental Economics*, Elsevier, pp. 325–354.
URL: <https://www.sciencedirect.com/science/article/pii/S1574009903010131>
- O’Brien, B. & Hagerty, N. (2025), ‘The trade-off between wildlife conservation and renewable energy: Evidence from eagles and wind turbines’, *Land Economics* **101**(3), 285–303. Preprint dated December 31, 2024 at <https://hagertynw.github.io/webfiles/Golden-Eagles-Versus-Wind-Turbines.pdf>.
URL: <https://doi.org/10.3368/le.101.3.092023-0091R>

- Office of Natural Resources Revenue (2025), ‘Federal revenue by location, fiscal years 2004–2025’, U.S. Department of the Interior, Natural Resources Revenue Data. Federal onshore revenue (royalties, bonuses, rents, other), FY2015, summed over AZ, CA, CO, ID, MT, NV, NM, OR, UT, WA, WY. Accessed September 25, 2026.
URL: https://revenuedata.doi.gov/downloads/fiscal_year_revenue.csv
- Parkhomenko, A. (2023), ‘Local causes and aggregate implications of land use regulation’, *Journal of Urban Economics* **138**, 103605.
URL: <https://www.sciencedirect.com/science/article/pii/S009411902300075X>
- Peltzman, S. (1976), ‘Toward a more general theory of regulation’, *The Journal of Law and Economics* **19**(2), 211–240.
URL: <https://doi.org/10.1086/466865>
- Prendergast, C. (2003), ‘The limits of bureaucratic efficiency’, *Journal of Political Economy* **111**(5), 929–958.
- Renard, K. G., Foster, G. R., Weesies, G. A., McCool, D. K. & Yoder, D. C. (1997), Predicting soil erosion by water: A guide to conservation planning with the revised universal soil loss equation (RUSLE), Technical Report Agriculture Handbook No. 703, U.S. Department of Agriculture, Agricultural Research Service, Washington, D.C.
- Reynaert, M., Souza-Rodrigues, E. & van Benthem, A. A. (2024), ‘The environmental impacts of protected area policy’, *Regional Science and Urban Economics* **107**, 103968. Urban Economics and the Environment.
URL: <https://www.sciencedirect.com/science/article/pii/S0166046223001035>
- Rimbey, N. R., Torell, L. A. & Tanaka, J. A. (2007), ‘Why grazing permits have economic value’, *Journal of Agricultural and Resource Economics* **32**(1), 20–40.
URL: <http://www.jstor.org/stable/40987349>
- Rio Puerco Field Office, New Mexico (2024), Rio puerco record of decision and approved resource management plan, Technical report, U.S. Department of the Interior, Bureau of Land Management, 100 Sun Avenue NE, Albuquerque, New Mexico, 87109.
URL: <https://eplanning.blm.gov/eplanning-ui/project/64954/510>
- Rock Springs Field Office, High Desert District (2024a), Rock springs proposed resource management plan and final environmental impact statement, Technical report, U.S. Department of the Interior, Bureau of Land Management, 280 Highway 191 North, Rock Springs, Wyoming, 82901.
URL: <https://eplanning.blm.gov/eplanning-ui/project/13853/510>
- Rock Springs Field Office, High Desert District (2024b), Rock springs record of decision and approved resource management plan, Technical report, U.S. Department of the Interior, Bureau of Land Management, 280 Highway 191 North, Rock Springs, Wyoming, 82901. See Management Actions 7002–7003, National Trail Management Corridor.
URL: <https://eplanning.blm.gov/eplanning-ui/project/13853/510>
- Rollet, V. (2025), Zoning and the dynamics of urban redevelopment, Technical report, Massachusetts Institute of Technology.
- Royal Gorge Field Office, Rocky Mountain District (2024), Eastern colorado record of decision and approved resource management plan, Technical report, U.S. Department of the Interior, Bureau of Land Management, 3028 East Main Street, C  non City, Colorado, 81212.
URL: <https://eplanning.blm.gov/eplanning-ui/project/68062/510>
- Scharfstein, D. S. & Stein, J. C. (1990), ‘Herd behavior and investment’, *American Economic Review* **80**(3), 465–479.

- Slattery, C. (2025), ‘Bidding for firms: Subsidy competition in the united states’, *Journal of Political Economy* **133**(8), 2563–2614.
URL: <https://www.journals.uchicago.edu/doi/abs/10.1086/735509>
- Steinbach, M. S. & Thomas, J. W. (2007), ‘Potential outcomes and consequence of a proposed grazing permit buyout program’, *Rangeland Ecology & Management* **60**(1), 36–44.
URL: <https://www.sciencedirect.com/science/article/pii/S1550742407500072>
- Stigler, G. J. (1971), ‘The theory of economic regulation’, *The Bell Journal of Economics and Management Science* **2**(1), 3–21.
URL: <http://www.jstor.org/stable/3003160>
- Stilwell, D. P. (2013), Final reasonable foreseeable development scenario for oil and gas, rock springs field office, wyoming, Reasonably foreseeable development scenario, U.S. Department of the Interior, Bureau of Land Management, Wyoming State Office, Reservoir Management Group.
- Train, K. E. (2009), *Discrete Choice Methods with Simulation*, 2 edn, Cambridge University Press, Cambridge.
- Uncompahgre Field Office (2019), Uncompahgre proposed resource management plan and final environmental impact statement, Technical Report DOI-BLM-CO-S050-2010-0001-RMP-EIS, U.S. Department of the Interior, Bureau of Land Management, Montrose, Colorado.
- Uncompahgre Field Office (2020), Record of decision and approved resource management plan, Technical Report BLM/CO/PL-20/009, U.S. Department of the Interior, Bureau of Land Management, Montrose, Colorado.
URL: <https://eplanning.blm.gov/eplanning-ui/project/64470/510>
- U.S. Department of Agriculture, Forest Service (1999), Roads analysis: Informing decisions about managing the national forest transportation system, Technical Report Miscellaneous Report FS-643, U.S. Department of Agriculture, Forest Service, Washington, D.C.
- U.S. Department of Agriculture, Natural Resources Conservation Service (2021), Title 190—rangeland hydrology and erosion handbook: The RHEM guide, Technical Report 190-647-H, 1st Edition, U.S. Department of Agriculture, Natural Resources Conservation Service, Washington, D.C.
- U.S. Fish and Wildlife Service (2018a), ‘Golden Eagle Resource Equivalency Analysis: Mortality’, Excel workbook, version dated October 22, 2018. Sheet “Debit Summary”: direct loss of 5.0655 present-value bird-years per eagle taken; sheet “Mitigation Owed”: 13.76 poles at 1:1 or 16.51 poles at 1.2:1 with 30-year retrofits; sheet “Bird-Years per Fledgling”: 3.8314 present-value bird-years per fledgling after first-year survival.
URL: <https://www.fws.gov/sites/default/files/documents/golden-eagle-mortality-rea.xlsx>
- U.S. Fish and Wildlife Service (2018b), ‘Golden Eagle Resource Equivalency Analysis: Nest Disturbance’, Excel workbook, version dated October 22, 2018. Sheet “Mitigation Owed”: one pair disturbed for one year is debited 0.59 fledglings, or 3.1448 present-value bird-years, and priced at 8.54 poles (\$64,066) at 1:1 or 10.25 poles (\$76,879) at the 1.2:1 ratio, with 30-year retrofits at \$7,500 per pole.
URL: <https://www.fws.gov/sites/default/files/documents/golden-eagle-nest-disturbance-rea.xlsx>
- U.S. Fish and Wildlife Service (2022), Protections for Raptors, Guidance, U.S. Fish and Wildlife Service, Wyoming Ecological Services Field Office, Cheyenne, Wyoming. Dated March 9, 2022. Table 1 lists the office’s recommended spatial and seasonal buffers for breeding raptors.
URL: <https://www.fws.gov/sites/default/files/documents/wyoming-ecological-services-field-office-raptor-guidelines-2022-03-09.pdf>

- U.S. Fish and Wildlife Service (2024), Final Environmental Assessment: 2024 Eagle Take Permit Rulemaking, Final environmental assessment, U.S. Fish and Wildlife Service, Migratory Birds and Habitat Program. Section 4.7.1.1, p. 73: the Service estimates that retrofitting 16 power poles, at \$7,500 apiece, offsets one golden eagle at a 1.2:1 ratio, and uses the resulting \$120,000 per eagle throughout the analysis.
URL: <https://www.fws.gov/sites/default/files/documents/2024-04/eagle-rule-final-ea-2024.pdf>
- Volcovici, V. & Renshaw, J. (2021), ‘New mexico asks white house for help weathering federal drilling pause’, Reuters, February 12, 2021.
- Voting and Election Science Team (2018), ‘2016 precinct-level election results’, Harvard Dataverse.
- Walls, M., Lee, P. & Ashenfarb, M. (2020), ‘National monuments and economic growth in the american west’, *Science Advances* **6**(12), eaay8523.
URL: <https://www.science.org/doi/abs/10.1126/sciadv.aay8523>
- Williams, H. F. L., Havens, D. L., Banks, K. E. & Wachal, D. J. (2008), ‘Field-based monitoring of sediment runoff from natural gas well sites in Denton County, Texas, USA’, *Environmental Geology* **55**(7), 1463–1471.
- Woods, N. D. (2013), ‘Regulatory democracy reconsidered: The policy impact of public participation requirements’, *Journal of Public Administration Research and Theory* **25**(2), 571–596.
URL: <https://doi.org/10.1093/jopart/mut042>

Figure 1: The Share Conserved vs. the Development Potential




Notes: Figure plots the share of each township^a conserved against the township's expected revenue from fluid leasable development.^b Each line is the relationship estimated by a logistic regression of the share conserved on log revenue, fit separately for each office and drawn over the middle 90% of that office's townships. The horizontal axis is development potential: the expected average annual revenue within a township, summed across fluid resources, deflated to 2014 dollars and plotted on a log scale. Townships with potential below \$100 a year (15% of the sample) cannot be placed on the log scale and are dropped. Blue curves fall as potential rises (8 of the 11 offices); gray curves are flat or rise (Buffalo, Wyoming, and Colorado River Valley and Uncompahgre, Colorado). Panel (a) shows three Wyoming offices: markers are the mean share and mean log potential within eight bins, each holding an equal number of the office's townships, and shaded areas are 95% confidence intervals for the fitted curve. Panel (b) shows every office's fitted curve without markers or intervals. The black curve pools the townships of all eleven offices, weighting each township equally, and is drawn over the 2nd to 98th percentiles of pooled potential; its shaded area is a 95% confidence interval with standard errors clustered by office.

^aTownships (6x6 mile squares) are standard land-survey units in the Western United States.


^bThe fluid leasable resources are conventional oil and gas, coalbed natural gas, and geothermal, whichever of the three the plan projects development for.

Figure 2: Considered Zoning Configurations around Two Natural Features




Notes: The nine maps show one area of the Four Rivers, Idaho planning area. The columns are three configurations the office considered: a development-oriented one (left), the chosen plan (middle), and a conservation-oriented one (right). Row (a) is every decision the plan makes there, which together protect four features: the Oregon National Historic Trail (solid line), a critical groundwater area (dashed line), elk winter range (the discretionary ground covering most of the left column), and sage-grouse habitat (the no-surface-disturbance ground at upper right in the left column). Rows (b) and (c) keep only the zoning tied to the trail and to the groundwater area; tile labels paraphrase the plan. Darker fills are more restrictive: *Discretionary* allows drilling with limits on where and when; *No surface disturbance* bars surface drilling but allows wells angled in from adjacent land (Closed in the model; Section 3.2); *Closed* removes the land from leasing. White is land left *Open*—in rows (b) and (c), open with respect to the feature shown; grey is land outside the BLM’s jurisdiction.

Figure 3: Distribution of Chosen Protections for Three Kinds of Feature



Notes: Figure plots the protections chosen for three types of features across all offices in the estimation sample. From right to left, the panels plot the slope grade above which drilling is restricted, the restricted distance from streams and rivers, and the restricted distance around raptor nests. The x-axis shows the chosen extent in each plan's original units, and the y-axis shows the share of protective decisions choosing that extent. Labels below the x-axis show how many offices chose each extent for at least one feature. Features not protected in the chosen plan are excluded from the figure. The displayed bars account for at least 99% of protective decisions in each panel.


Figure 4: Willingness to Pay for Conservation, by Field Office and Feature Group



Willingness to pay for conservation (2014 dollars)


Notes: Figure plots each field office's estimated willingness to pay for conservation, in dollars of forgone revenue, from Section 2's model and Section 4.2's specification. Each bar is the estimate for one office and one natural-capital class; the whiskers are 90% confidence intervals from the Hessian standard errors. The horizontal axis is signed-log, matching Figure A3. "no data" marks an office×class group with no data.

Figure 5: Willingness to Pay for Conservation and the 2016 Trump Vote Share, by Feature Group



Notes: Figure plots each field office's estimated willingness to pay for conservation (Figure 4) against the share of the 2016 two-party presidential vote cast for Donald Trump in the office's planning area, with one panel for each non-market resource. Each marker is one office, and its shape denotes the office's state. The panels include the twenty-eight estimates that are statistically distinguishable from zero. The vote share sums precinct returns over the planning area; a precinct that straddles the area's boundary contributes in proportion to the share of its 2010 population living inside. The vote share data is from Voting and Election Science Team (2018). Lines are least-squares fits of log willingness to pay on the vote share for each resource type: the dark line leaves out Rio Puerco, New Mexico (hollow marker), and the light line includes it. Rio Puerco is set apart because it is the only New Mexico office, and New Mexico's voters lean Democratic while its state government depends heavily on revenue from federal onshore mineral leasing (New Mexico Legislative Finance Committee 2023, Volcovici & Renshaw 2021). Rio Puerco has no wildlife estimate, so the wildlife panel has one line. The vertical axis is on a log scale.

Figure 6: Protecting the Same Number of Birds with a Revised Standard or Price-Like Instruments



Notes: Figure plots changes in baby birds (fledglings) protected and revenue forgone under the standard revision and price-like instrument counterfactuals, relative to the baseline. Dark orange bars show the standard revision under *heterogeneous adoption*, in which all offices respond to a common one-mile recommendation but retain office- and nest-specific unobserved heterogeneity in the conservation standards they ultimately follow. Light and dark blue bars show two implementations of the price-like instrument, both calibrated to protect the same aggregate number of additional fledglings as the standard revision. The *minimum value* implementation (light blue) raises each office's value of avoiding wildlife disturbance to at least \$212,000 per relevance-weighted acre, while the *common value* implementation (dark blue) sets every office's value to \$240,000. The top panel reports aggregate changes, while the bottom panel breaks them out by office. The first column shows changes in fledglings protected, the second changes in revenue forgone, and the third, appearing only in the top panel, revenue forgone per additional fledgling. The "Other offices" category includes Uncompahgre, Colorado River Valley, Royal Gorge, and Grand Junction in Colorado, and Four Rivers in Idaho.

Table 1: Willingness to Pay and Cost of Departure, by Field Office

Field office	Willingness to pay (\$/wtd. acre)			Closure FE (logits)	Scale (log σ)		Cost of departure (κ)				
	Soil	Water	Wildlife		Any closed label	Soil+Water	Wildlife	Soil+Water		Wildlife	
								\$ per unit departure	at the mean departure	\$ per unit departure	at the mean departure
Rock Springs, WY	\$58,534*** (\$7,508)	\$53,944*** (\$8,065)	\$38,263*** (\$2,513)	−7.64*** (1.29)	3.73 (0.223)	1.34 (0.336)	\$2,408,524 (\$652,793)	\$299,692 <i>at 0.353</i>	\$353,696 (\$125,802)	\$10,918 <i>at 0.176</i>	
Casper, WY	\$55,185*** (\$3,783)	\$17,412*** (\$2,348)	\$26,966*** (\$2,158)	−1.92*** (0.0783)	3.45 (0.113)	2.91 (0.228)	\$1,960,468 (\$299,093)	\$52,143 <i>at 0.163</i>	\$1,948,109 (\$466,637)	\$86,143 <i>at 0.210</i>	
Buffalo, WY	\$31,858*** (\$3,625)	\$46,315*** (\$5,412)	\$321,292** (\$125,336)	−1.75*** (0.0582)	4.12 (0.169)	6.18 (0.496)	\$7,968,144 (\$1,621,067)	\$568,432 <i>at 0.267</i>	\$28,681,361 (\$14,305,417)	\$2,120,810 <i>at 0.272</i>	
Lander, WY	\$149,074*** (\$9,200)	\$82,070*** (\$10,380)	\$65,302 (\$66,252)	−0.385*** (0.104)	4.28 (0.130)	6.54 (1.07)	\$3,315,653 (\$590,033)	\$130,460 <i>at 0.198</i>	\$833,198,799 (\$909,503,317)	\$4,285,935 <i>at 0.072</i>	
Grand Jct., CO	\$162,631*** (\$31,181)	\$758,014*** (\$173,148)	\$3,630 (\$26,853)	1.46*** (0.164)	5.55 (0.234)	0.584 (0.813)	\$37,995,229 (\$9,634,168)	\$2,059,150 <i>at 0.233</i>	\$105,624 (\$93,117)	\$6,450 <i>at 0.247</i>	
Uncompahgre, CO	\$22,983*** (\$4,110)	\$112,826*** (\$19,503)	\$36,994*** (\$3,395)	0.755*** (0.119)	3.21 (0.257)	0.351 (0.343)	\$2,851,967 (\$946,023)	\$231,246 <i>at 0.285</i>	\$101,440 (\$37,081)	\$6,174 <i>at 0.247</i>	
Royal Gorge, CO	\$406,979*** (\$30,271)	\$1,605,124*** (\$97,120)	\$897,196*** (\$68,596)	1.64*** (0.0622)	2.80 (0.0840)	2.39 (0.385)	\$235,181,741 (\$22,437,427)	\$38,065 <i>at 0.013</i>	\$1,766,854 (\$699,877)	\$53,364 <i>at 0.174</i>	
Colo. River Valley, CO	\$71,295*** (\$2,054)	\$66,425*** (\$3,614)	\$95,551*** (\$4,167)	1.93*** (0.149)	3.83 (0.213)	1.29 (0.481)	\$2,207,466 (\$523,061)	\$331,449 <i>at 0.387</i>	\$7,869,325 (\$3,920,763)	\$25,286 <i>at 0.057</i>	
Little Snake, CO	−\$3,990 (\$6,907)	\$53,275*** (\$8,891)	\$27,964*** (\$4,510)	1.36*** (0.149)	4.16 (0.255)	−0.116 (0.357)	\$23,234,278 (\$6,358,336)	\$1,070,296 <i>at 0.215</i>	\$3,964,308 (\$1,496,850)	\$13,133 <i>at 0.058</i>	
Four Rivers, ID	\$28,157** (\$11,741)	\$299,305*** (\$56,004)	\$793 (\$4,824)	0.805*** (0.121)	0.785 (0.233)	−5.28 (0.547)	\$477,808 (\$115,284)	\$31,266 <i>at 0.256</i>	\$0 (≤ \$36)	\$0 <i>at 0.558</i>	
Rio Puerco, NM	\$84,054*** (\$19,744)	\$12,038*** (\$4,233)	<i>no data</i>	−1.05*** (0.138)	1.39 (0.270)	<i>no data</i>	\$104,651,030 (\$31,745,327)	\$3,750 <i>at 0.006</i>	<i>no data</i>	<i>no data</i>	

Notes: Table reports the estimates from Section 2's model and Section 4.2's specification, fit once on all eleven offices together (12,509 choice situations, log-likelihood −23,208.90). Willingness to pay is the revenue, in 2014 dollars, an office forgoes to protect an acre at the feature from development, for each natural-capital class. The closure effect is the office's residual tendency to close a feature outright rather than leave it under discretionary stipulations, common to all natural-capital classes and measured in log-odds rather than dollars. The scale, reported in logs for soil and water together and for wildlife, is the spread of the idiosyncratic differences among a feature's alternatives. The cost of departure is reported in dollars for soil and water together and for wildlife; it is estimated in logs, so it is positive by construction and carries no significance stars. The *at the mean departure* columns report the penalty an office pays on an average protective decision: the cost of departure times the mean squared gap between its chosen buffer and the standard it follows, the proxy plus its correction, with the gap averaged over the office's latent types by their posterior probabilities given the choice. Open choices carry no departure and are left out. The italic line reads that departure as a distance: how far, on the relevance scale, the office's typical protective decision lands from the buffer its standard called for, where 0 means it stopped exactly at the standard and 0.3 means the last acre it protected mattered about a third less, or more, than its standard would have. *no data* marks a class×office group that I do not currently have data for. A cost printed as \$0 (wildlife at Four Rivers, ID) sits at the boundary of what the fit allows: the office is estimated to pay nothing to depart there, and the parenthesis is the one-sided 95% upper bound. Standard errors in parentheses, by the delta method for the cost of departure; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 2: How Often an Adjacent Buffer’s Net Benefits Exceed the Predicted Buffer’s

Field office	Decisions	Share of decisions with a net benefit gain			
		> \$0	> \$100	> \$1k	> \$10k
Rock Springs, WY	1,127	83%	64%	26%	6%
Casper, WY	1,902	67%	49%	33%	12%
Buffalo, WY	2,696	94%	93%	89%	77%
Lander, WY	518	73%	70%	58%	31%
Grand Junction, CO	221	50%	48%	42%	29%
Uncompahgre, CO	386	51%	39%	25%	9%
Royal Gorge, CO	1,427	44%	38%	27%	11%
Colorado River Valley, CO	454	69%	53%	32%	22%
Little Snake, CO	427	89%	87%	84%	64%
Four Rivers, ID	345	47%	18%	8%	3%
Rio Puerco, NM	195	72%	18%	9%	3%
Median office		69%	49%	32%	12%

Notes: Table reports, for the decisions in which the fitted model predicts protection, the share in which a buffer adjacent to the predicted buffer has a higher net benefit—its conservation value net of the revenue given up—by the size of the gain. The predicted buffer is the protective alternative with the highest predicted probability; its adjacent buffers are the alternatives under the same regime one step narrower and one step wider. Figure A7 plots the first share column by natural-capital class. The Decisions column documents the number of decisions for which the model predicts protection. The first share column counts a gain of any amount; the others count only gains above \$100, \$1,000 and \$10,000 a year (2014 dollars). The last row is the median across offices of each column.

A Economic Interpretation of Land Use Plan Regulation

The Federal Land Policy and Management Act (FLPMA) serves two central functions in BLM land-use planning. First, it grants the agency comprehensive authority to manage public lands under a mandate of multiple use and sustained yield. This mandate requires the BLM to balance competing uses—such as grazing, energy development, recreation, and conservation—while ensuring that public benefits are sustained over time. Second, FLPMA specifies the procedural mechanism through which the agency is to pursue this mandate: Resource Management Plans (RMPs), commonly referred to as land-use plans (LUPs).

Under the National Environmental Policy Act (NEPA), LUP development must include all required components of an Environmental Impact Statement (EIS) (Bureau of Land Management 2008). This process requires the BLM to evaluate the environmental and economic impacts of their choices, assess and disclose alternative management strategies, and engage with the public and other jurisdictions throughout plan development (Muhn & Stuart 1988).

Together, NEPA and FLPMA require the BLM to publicly define and solve an optimization problem through land-use planning. FLPMA defines the objective—management under a multiple-use, sustained-yield mandate—while NEPA requires the agency to articulate the feasible set by evaluating a full range of reasonable land-use alternatives. Both statutes require planners to publicly identify the market and non-market resource endowments that enter the objective function. The adopted LUP represents the chosen solution, specifying how each resource is governed over the plan’s tenure. This paper’s goal is to recover the objective function.

The remainder of this section shows how regulations, implementing NEPA and FLPMA, define the objective and informational content of BLM land-use plans. I also link this regulatory language to the paper’s analytical framework.

A.1 The Objective of Land Use Plans

This section outlines the objective of LUPs as stated in regulation.

LUPs’ objective is the same as the BLM’s mandate: “maximize resource values for the public through a rational, consistently, applied set of regulations and procedures which promote the concept of multiple use management” (CFR 43 Part 1601.0-2).⁴⁸ In this context, “multiple use management” is defined in CFR 43 Part 1601.0-5. The following excerpts highlight key elements of that definition, though they do not represent it in full:

1. “making the most judicious use of the lands for some or all of these resources or related services over areas large enough to provide sufficient latitude for periodic adjustments in use to conform to changing needs and conditions; the use of some lands for less than all of the resources”
2. “harmonious and coordinated management of the various resources...with consideration being given to the relative values of the resources and not necessarily to the combination of uses that will give the greatest economic return or the greatest unit output”

To map this legal definition of the LUP objective into a mathematical analogue, I summarize the mandate as follows: land-use plans should maximize resource values, accounting for both market and non-market values, while balancing competing uses. In economic terms, offices choose allocations to maximize a weighted sum of resource endowments, where the weights represent the relative values assigned to each resource.

A.2 Feasible Solution Set

In this section, I outline the federal regulations requiring that land-use plan development specify not only a final zoning configuration, but also a set of alternative configurations that span the space of feasible options—the considered configurations, in this paper’s language—thereby defining the feasible choice set.

⁴⁸See Section B for the full subpart.

43 CFR Part 1600 requires that the LUP development process “consider all reasonable resource management alternatives and develop several complete alternatives for detailed study” (CFR 43 1610.4-5). The regulation’s “alternatives” are zoning configurations in the vocabulary of this paper, and those developed for detailed study are its considered configurations. This legal requirement defines the *full* set of alternatives which address a LUP’s goal through different means and are *feasible*.

In economic terms, I interpret the regulatory requirement to consider “reasonable” alternatives, as a requirement to identify *feasible* solutions to the land use planning optimization problem. Specifically, the regulation defines “reasonable” alternatives as those that:

- Respond to the LUP’s “purpose and need” and address the “issues” identified during the LUP scoping process.⁴⁹
- Are feasible given estimated costs, logistics, and available technology.
- Contain internally consistent components; for example, an alternative should not both prohibit visual disturbances and allow oil and gas drilling on the same parcel.

I interpret the regulatory requirement to consider “all” alternatives as a requirement to define alternatives which span the feasible set of solutions. This interpretation is supported by guidance in how the BLM planning process should apply NEPA (Bureau of Land Management 2008):

“For some proposals there may exist a very large or even an infinite number of possible reasonable alternatives. When there are potentially a very large number of alternatives, you must analyze only a reasonable number to cover the full spectrum of alternatives (see Question 1b, CEQ, Forty Most Asked Questions Concerning CEQ’s NEPA Regulations, March 23, 1981).”

A.3 Resource Endowments

In this section, I outline the federal regulations requiring land-use plan developers to evaluate the impacts of their decisions, thereby documenting their beliefs about resource endowments across the planning area. I also provide an example of this analysis from a land-use plan.

For each zoning alternative, including the final plan—each considered configuration and the adopted one, in this paper’s language—planners must publicly outline the predicted impact of the given alternative on physical, biological, economic and social resources in the planning area (CFR 43 Part 1610.4-6). This assessment begins with a detailed description of existing resource endowments and then evaluates how each zoning alternative would affect those resources, identifying both positive and negative impacts. This second step closely resembles a benefit–cost analysis, in which benefits and costs are quantified in terms of the acreage of each resource that would be affected by a given zoning choice. For example, in the Four Rivers Field Office Draft RMP Amendment, the planners quantify the impact of alternate zoning configurations on vegetation resources in terms of the acres exposed to “surface disturbing activities”. Specifically, the planning documents write:

“The most acreage of right-of-way (ROW) exclusion, closure or open with a no-surface-occupancy (NSO) stipulation to fluid leasable minerals, and renewable energy exclusion areas would occur under Alternative B, which would reduce surface disturbing impacts from development on vegetation and help maintain intact vegetation communities in those areas. The least amount of ROW exclusion would occur under Alternative C, resulting in the greatest potential impacts from ROW development on vegetation.”

The documentation then goes on to detail acreage of permitted surface disturbance by vegetation community under the different configurations considered. Figure A2 shows a table from the planning documents summarizing this analysis at the planning area level. In this example, each row of the table details the

⁴⁹ An LUP’s “purpose and need” section is a required component that outlines the rationale for the plan. See 40 CFR §1502.13 for the relevant regulatory requirement. In addition, under 43 CFR §1610.4-1 and 40 CFR §1501.7, the scoping process must identify “issues” to be addressed in the LUP. These issues must be within the scope of analysis, describe a relationship between an action and a resource, be scientifically analyzable, and must not be resolved solely by existing law, regulation, or policy.

acres of a given vegetation community where ground disturbance would be allowed under each zoning alternative, labeled as “Alt. A”, “Alt. B”, “Alt. C” and “Alt. D”. Earlier and later discussions in the planning documents highlight that the vegetation communities: Evergreen Forest, Shrub-Steppe & Perennial Grassland, and Riparian & Wetland are valuable vegetation communities, either for logging, preservation of protected species, wildfire resilience, etc. While the vegetation community “Exotic Annual Grassland” have negative value, reducing habitat for native species and decreasing wildfire resilience.

Mineral and energy resources requirements for data collection and analysis follow a more structured protocol. Bureau of Land Management (1985) requires each plan to rank the planning area’s occurrence potential and development potential on a categorical scale (none, low, moderate, high), where development potential combines the geological ranking with historic drilling and recent industry interest to state where the office expects development within the life of the plan. Appendix F describes how these development-potential categories, together with the plan’s development forecast, yield the per-acre revenue and disturbance densities of Section 3.3.

B Relevant BLM Land Use Planning Regulation

B.1 CFR 43 Part § 1601.0-2 Objective

The objective of resource management planning by the Bureau of Land Management is to maximize resource values for the public through a rational, consistently applied set of regulations and procedures which promote the concept of multiple use management and ensure participation by the public, state and local governments, Indian tribes and appropriate Federal agencies. Resource management plans are designed to guide and control future management actions and the development of subsequent, more detailed and limited scope plans for resources and uses.

B.2 CFR 43 Part § 1601.0-5 Definitions

Multiple use means the management of the public lands and their various resource values so that they are utilized in the combination that will best meet the present and future needs of the American people; making the most judicious use of the lands for some or all of these resources or related services over areas large enough to provide sufficient latitude for periodic adjustments in use to conform to changing needs and conditions; the use of some lands for less than all of the resources; a combination of balanced and diverse resource uses that takes into account the long term needs of future generations for renewable and non-renewable resources, including, but not limited to, recreation, range, timber, minerals, watershed, wildlife and fish, and natural scenic, scientific and historical values; and harmonious and coordinated management of the various resources without permanent impairment of the productivity of the lands and the quality of the environment with consideration being given to the relative values of the resources and not necessarily to the combination of uses that will give the greatest economic return or the greatest unit output.

B.3 CFR 43 Part § 1601.0-6 Environmental impact statement policy

Approval of a resource management plan is considered a major Federal action significantly affecting the quality of the human environment. The environmental analysis of alternatives and the proposed plan shall be accomplished as part of the resource management planning process and, wherever possible, the proposed plan and related environmental impact statement shall be published in a single document.

B.4 43 CFR Part § 1610.4-5 Formulation of Alternatives

“At the direction of the Field Manager, in collaboration with any cooperating agencies, BLM will consider all reasonable resource management alternatives and develop several complete alternatives for detailed study. Nonetheless, the decision to designate alternatives for further development and analysis remains the exclusive responsibility of the BLM. The alternatives developed shall reflect the variety of issues and guidance applicable to the resource uses. In order to limit the total number of alternatives analyzed in detail to a manageable number for presentation and analysis, all reasonable variations shall be treated as sub-alternatives. One alternative shall be for no action, which means continuation of present level or systems of resource use. The plan shall note any alternatives identified and eliminated from detailed study and shall briefly discuss the reasons for their elimination.”

NEPA Section 102(2)(E) directs the BLM to “study, develop, and describe appropriate alternatives to recommended courses of action in any proposal that involves unresolved conflicts concerning alternative uses of available resources” (Bureau of Land Management 2008). Specifically, the BLM handbook on NEPA compliance states that (Bureau of Land Management 2008):

“The range of alternatives [should explore] alternative means of meeting the purpose and need for the action. ...the purpose and need statement⁵⁰ helps define the range of alternatives. The broader the purpose and need statement, the broader the range of alternatives that must be analyzed.⁵¹ You must analyze those alternatives necessary to permit a reasoned choice (40 CFR 1502.14). For some proposals there may exist a very large or even an infinite number of possible reasonable alternatives. When there are potentially a very large number of alternatives, you must analyze only a reasonable number to cover the full spectrum of alternatives (see Question 1b, CEQ, Forty Most Asked Questions Concerning CEQ’s NEPA Regulations, March 23, 1981). When working with cooperating agencies, your range of alternatives may need to reflect the decision space and authority of other agencies, if decisions are being made by more than one agency.”

B.5 CFR 43 Part § 1610.4-6 Estimation of effects of alternatives

The Field Manager, in collaboration with any cooperating agencies, will estimate and display the physical, biological, economic, and social effects of implementing each alternative considered in detail. The estimation of effects shall be guided by the planning criteria and procedures implementing the National Environmental Policy Act. The estimate may be stated in terms of probable ranges where effects cannot be precisely determined.

C Examples from the Planning Record

This section provides examples from the planning record supporting assertions with regards to how allocation decisions are made. Each subsection names an assertion made in the main text and gives the examples that support it. All examples are from planning documents published by one of the eleven offices in the estimation sample.

C.1 Choices Not Fixed by Ecology, Hydrology, or Geology

This section provides examples supporting the assertion that an allocation decision, the restriction an office places on a feature and the distance it extends, is a choice rather than something the feature’s ecology, hydrology, or geology determines.

⁵⁰The purpose and need statement of an RMP or RMP amendment “shall briefly specify the underlying purpose and need to which the agency is responding in proposing the alternatives including the proposed action” (40 CFR 1502.13); the BLM NEPA handbook quotes the same language in its §6.2 (Bureau of Land Management 2008).

⁵¹Alternatives are only “reasonable” if they respond to the purpose and need for the RMP or RMP amendment. Only reasonable alternatives need to be considered (Bureau of Land Management 2008).

Grand Junction, Colorado. Colorado Parks and Wildlife asked the Grand Junction office to raise its 30 m buffer on ephemeral streams to 100 m. Existing literature, already cited in the plan itself, places that 30 m at the low end of the range of effective riparian buffers (Johnson & Buffler 2008, p. 26). The office declined on the ground that “the minimum buffer of 30 meters will allow flexibility for protecting these systems according to site-specific circumstances” (Grand Junction Field Office 2015, p. 6-19).

Uncompahgre, Colorado. The office’s existing plan prohibits surface use within a half mile of an osprey nest, on the ground that “[t]he sensitivity of osprey to human associated disturbance activities requires a half-mile buffer zone to avoid nest abandonment” (Uncompahgre Field Office 2019, p. B-105). The approved plan cuts the radius to a quarter mile, offering no rationale for the revision and no reassessment of the sensitivity the earlier stipulation invoked (Uncompahgre Field Office 2020, p. B-33).

Casper, Wyoming. Casper’s approved plan is explicit that development on slopes over 25 percent is hazardous, yet the office does not close those slopes to surface disturbance outright. It prohibits surface disturbance “unless individual site plans are submitted to and approved by the authorized officer” (Casper Field Office, Wyoming State Office 2007, Decision 1022).

C.2 Decisions Made Feature by Feature Rather Than by Feature Type

This section provides examples of offices zoning one feature differently from the rest of its type.

Rock Springs, Wyoming. Rock Springs closes the saline soils of the Red Creek watershed to fluid mineral leasing while comparably saline soils elsewhere in the planning area remain open to leasing; the alternatives considered would have left Red Creek open as well (Rock Springs Field Office, High Desert District 2024*b*). The office likewise closes the aquifer recharge areas around the towns of Superior and McKinnon to drilling but assigns discretionary restrictions to other mapped recharge areas (Rock Springs Field Office, High Desert District 2024*b*).

Rio Puerco, New Mexico. One contiguous portion of the planning area’s winter elk habitat is Closed while the rest receives discretionary protection, a seasonal limit on drilling (Rio Puerco Field Office, New Mexico 2024).

C.3 Different Decisions for Different Market Resources at the Same Feature

This section provides examples of offices restricting different market resources differently at the same feature.

Four Rivers, Idaho. The Four Rivers plan amendment separates oil and gas from geothermal management actions, closing areas with low or no oil and gas potential while leaving geothermal management unchanged (Bureau of Land Management 2023).

Rock Springs, Wyoming. Rock Springs restricts oil shale within a quarter mile of a national historic trail while restricting all other resources within five miles of it (Rock Springs Field Office, High Desert District 2024*b*).

C.4 Closures Opposed Whether or Not Development Was Expected

This section provides examples of closure being opposed as a decision in itself, whether or not development was expected on the land closed.

Rock Springs, Wyoming. A trona operator protesting the Rock Springs draft plan objected to the closure of solid mineral leasing on slopes greater than 25%, which under that alternative covered 2,119,920 acres—of which 49,224, or 2.3 percent, fall inside the one area the plan identifies as materially affecting trona development (Rock Springs Field Office, High Desert District 2024a, ES-6, 4-84, W-20–W-21). The operator asked why BLM was “prohibiting solid mineral leasing in areas where soils have . . . slopes greater than 25%,” when “most of the RMP area exhibits at least one of these characteristics, as the area is a high desert” (Rock Springs Field Office, High Desert District 2024a, W-20–W-21).

C.5 Disturbance Matters Less the Farther It Is from the Feature

This section provides an example of an office treating disturbance as mattering less the farther it lies from the feature.

Lander, Wyoming. Lander’s proposed plan protects riparian areas with a 500-foot buffer. One alternative the office considered would have protected more. The office judged that 500 feet “would likely be adequate” to provide “the necessary attributes for water dissipation, sediment capture, and habitat needs,” and that “it is currently undetermined what level of beneficial impact there would be from the additional buffer” under the wider alternative (Lander Field Office, Wind River/Bighorn Basin District 2013, Ch. 4, p. 833).

C.6 Decisions Regarding Prairie Dog Towns Factor in Their Size

This section provides examples of offices factoring prairie dog town size when making protection town protection decisions.

Little Snake, Colorado. Little Snake prohibits disturbance of more than one acre in active white-tailed prairie dog towns smaller than ten acres and permits it in larger ones (Little Snake Field Office 2010, p. Q-237). When a commenter objected that the rule “would conserve relatively small colonies while allowing for development in larger colonies even though these larger colonies have greater biological value” and “promotes habitat fragmentation.” The office justified their choice by arguing it “would preserve smaller colonies that are more sensitive to surface disturbance while allowing for some disturbance within larger colonies that could more easily adapt to, recover from, or absorb impacts associated with surface disturbances” (Little Snake Field Office 2010, p. Q-67).

Uncompahgre, Colorado. Uncompahgre considered two designs for prairie dog towns, each in a different alternative. One protects every active town the same way, with a 150-foot ring around it in which surface use is prohibited. The other protects towns by size: in towns smaller than ten acres no activity may disturb more than one acre, while larger towns carry no such limit (Uncompahgre Field Office 2019, p. B-38).

Lander, Wyoming. Lander weighed prohibiting disturbance in all colonies against avoiding it only in “complexes larger than 100 acres” (Lander Field Office, Wind River/Bighorn Basin District 2013, Ch. 2, p. 131).

D Standards in the Planning Record

This section provides examples from the planning record supporting assertions with regards to how allocation decisions are publicly justified. The first subsection sets out the legal requirements and agency guidance that oblige an office to justify each restriction in writing when it responds to public comment. The remaining subsections give examples of a common way of meeting that obligation, citing a standard set elsewhere and stating that the choice is consistent with it, as well as of the departures from standards that offices defend on their own and of the limits on which standards an office will cite. Each subsection names an assertion made in the main text and gives the examples that support it. All examples are from planning documents

published by one of the eleven offices in the estimation sample or from the agency guidance that governs them.

D.1 The Public Comment Process

This section sets out the regulatory requirements and agency guidance that underpin the formal public review every land-use plan goes through. Regulation requires a written response to every substantive comment, followed by a protest review and the Governor’s consistency review. The agency’s handbook directs an office to adopt the least restrictive measure that meets a resource objective and to justify each restriction in the plan. The commenters press in opposite directions, and the office owes each of them a written answer.

Regulatory requirements. A plan is published in draft for public comment and again, as a proposed plan, for protest, before a final decision is made. The regulation in force when most of the sample’s plans were written says the agency “shall assess and consider comments both individually and collectively, and shall respond” in the final document. One permitted response is to explain why a comment warrants no change, “citing the sources, authorities, or reasons which support the agency’s position” (40 CFR 1503.4). The BLM’s NEPA handbook states the duty without qualification: “[w]hen preparing a final EIS, you must respond to all substantive written comments” (Bureau of Land Management 2008, p. 65). Once the office publishes its proposed plan, any person who participated in the planning may *protest* it. The Director must then rule on the protest in writing, either upholding the office’s decision with reasons or requiring the office to change the plan before it is approved (43 CFR 1610.5-2).

BLM’s recommendation for justifying decisions. For fluid mineral leasing, agency guidance recommends offices justify a restriction by showing that less restrictive measures were considered and found inadequate, and that they state that justification in the plan. Specifically, the BLM’s handbook directs that “[t]he least restrictive stipulation that effectively accomplishes the resource objectives or uses for a given alternative should be used” (Bureau of Land Management 2013, p. III-11). It asks the plan to show that a milder option was weighed: “[i]f a closure or operating constraint is discretionary with the BLM, evidence that a less restrictive mitigation measure was considered should be reflected in the range of alternatives analyzed in detail” (Bureau of Land Management 2013, p. III-13). For the alternative the office prefers, the plan “should provide evidence that less restrictive measures were considered but found inadequate” (Bureau of Land Management 2013, p. III-14). The plan’s text “should summarize the justification for constraints, nondiscretionary and discretionary closures, including references to any applicable laws, executive orders, etc.” (Bureau of Land Management 2013, p. IV-1). Grand Junction, Colorado follows this instruction literally. It tells its commenters that “[t]he justification for each stipulation can be found in Appendix B under the specific stipulation” (Grand Junction Field Office 2015, p. 6-280), and each stipulation there ends with a line headed “JUSTIFICATION” (Grand Junction Field Office 2015, pp. B-40–B-41).

Who comments. Commenters on the same restriction frequently ask for opposite changes. On a single page of Grand Junction’s response to comments, the county commissioners say the office “has not justified the need” for the buffer it proposed, and the state wildlife agency asks for a stricter regime (Grand Junction Field Office 2015, p. 6-45). Elsewhere in the same volume an industry group asks for the “scientific justification” of a river buffer, and a conservation group calls the stream buffers “not adequate” (Grand Junction Field Office 2015, pp. 6-270, 6-276). Little Snake, Colorado’s response table carries comments from a gas operator, the county commissioners, two conservation groups, and the U.S. Fish and Wildlife Service (Little Snake Field Office 2010, pp. Q-236–Q-247, Q-377–Q-378). In these records operators, counties, and industry groups generally press for less protection, and conservation groups and wildlife agencies generally press for more. The office owes each of them a written answer.

Comments that changed a plan: Buffalo and Little Snake. Buffalo’s draft drew comments in both directions on its raptor buffers, and the office revised the raptor actions “to fully conform to the distances

and dates recommended by the USFWS Wyoming Ecological Services Office” (Buffalo Field Office, High Plains District 2015*b*, p. 30). Little Snake’s draft carried forward an eighth-of-a-mile buffer around raptor nests from its earlier plan; a conservation group asked for a quarter mile, and the approved plan adopted it, finding the increase “necessary to protect the raptor species, as well as to be consistent with state-wide stipulations throughout Colorado” (Little Snake Field Office 2011, p. ROD-2). Both exchanges are given in full in Sections D.4 and D.5.

D.2 Standards That Predate the Plan

This section provides examples supporting the assertion that when offices set a buffer, there are already published distances it can adopt instead of defending one of its own.

Wyoming: the Mitigation Guidelines. The BLM in Wyoming maintains a set of *Mitigation Guidelines for Surface-Disturbing and Disruptive Activities*, updated “for all Field Offices in Wyoming.” Subject to exceptions, they prohibit surface disturbance on “[s]lopes in excess of 25 percent,” “[w]ithin 500 feet of surface water and/or riparian areas,” and “[w]ithin either one-quarter mile or the visual horizon (whichever is closer) of historic trails” (Casper Field Office, Wyoming State Office 2007, Appendix I, p. I-2). The guidelines state their own purpose: they “are primarily for the purpose of attaining statewide consistency in how requirements are determined for avoiding and mitigating environmental impacts” (Casper Field Office, Wyoming State Office 2007, Appendix I, p. I-1). They predate any one plan. “All BLM RMPs have included these guidelines as appendices,” and “[p]ublic comment on the guidelines, per se, has not been requested” (Casper Field Office, Wyoming State Office 2007, Appendix I, p. I-1). They are also the reference point for the alternatives a plan considers. Planners use them “to develop a baseline for measuring and comparing impacts among the alternatives” and “to help determine whether more stringent or less stringent mitigations should be considered” (Casper Field Office, Wyoming State Office 2007, Appendix I, p. I-1). Lander, Wyoming’s plan reprints the same text (Lander Field Office, Wind River/Bighorn Basin District 2013, pp. 1595–1596). The national handbook, written in 1990, already takes the “Wyoming BLM Standard Stipulations for Surface Disturbing Activities” as its worked example of how to describe an alternative (Bureau of Land Management 2013, Illustration 3, p. 1).

Wildlife agencies. Colorado’s state wildlife agency published *Recommended Buffer Zones and Seasonal Restrictions for Colorado Raptors* in 2008. Grand Junction’s plan, approved in 2015, names that document as the justification for each of its raptor stipulations. The half-mile buffer around ferruginous hawk nests, for example, “is necessary to protect ferruginous hawk nesting habitat per” the state’s recommended buffer zones (Grand Junction Field Office 2015, p. B-40). The U.S. Fish and Wildlife Service’s Wyoming office publishes its own table of recommended raptor buffers and dates. For special status raptors, Buffalo’s proposed plan prints no distances of its own. It adopts the Service’s “recommended spatial buffers and dates for breeding raptors” by reference to the Service’s web page (Buffalo Field Office, High Plains District 2015*b*, pp. 217–220).

Statewide stipulations. The BLM’s state offices issue stipulations meant to be used by every field office in the state. Grand Junction describes the Colorado effort: “[p]art of BLM leasing reform is to enact minimum stipulations (for the various programs) that could be applied across the state,” and the group working on water “reduced the number of stipulations from approximately 70 to approximately 12 in Colorado” (Grand Junction Field Office 2015, p. 6-271). Uncompahgre refers to “BLM statewide stipulations, which are used on BLM-administered lands across Colorado” (Uncompahgre Field Office 2019, p. R-223). The state wildlife agency supplied its own statewide list. It told Uncompahgre that in 2010 it had given the BLM’s State Office “recommendations for fluid mineral lease stipulations relevant to Resource Management Plan (RMP) revisions” (Uncompahgre Field Office 2019, p. R-209).

The office’s own earlier plan. Each office also inherits the plan it is revising. The handbook requires that “[e]xisting fluid minerals management, including constraints, is summarized and presented in the RMP/EIS as the ‘no action’ alternative,” which “serves as the baseline for discussing the other alternatives” (Bureau of Land Management 2013, p. III-13). The existing buffers therefore sit in every revision as a named alternative. In Rock Springs, Wyoming’s revision that alternative “represents a continuation” of the plans then in force, the older of them adopted in 1997 (Rock Springs Field Office, High Desert District 2024*b*, p. 1-3). Little Snake tells a commenter that “[t]he language in Alternative A was obtained from the 1989 Little Snake RMP and ROD and is not subject to revision within Alternative A” (Little Snake Field Office 2010, p. Q-377).

D.3 Commitments to Consistency Made Before Choosing

This section provides examples supporting the assertion that offices commit themselves to consistency with standards and with neighboring offices before they draft allocation decisions. Rock Spring, WY’s decision record explains that the approved stipulations “were proposed to align with similar restrictions in adjacent planning area field offices and to provide clarity and resolve potential conflicts from the Draft” (Rock Springs Field Office, High Desert District 2024*b*, pp. 1-4–1-5). Little Snake states, “Lease stipulations will be reviewed for consistency with neighboring BLM Field Offices and States, and where there are discrepancies efforts will be undertaken to achieve consistency or to explain any differences” (Little Snake Field Office 2011, pp. RMP-6–RMP-7). Colorado River Valley, Colorado states, “Lease stipulations will be reviewed for consistency with neighboring field offices and states. Where there are discrepancies, efforts will be made to achieve consistency” (Colorado River Valley Field Office 2014, p. 1-14). Both offices also commit to a common format for stipulations that a regional committee prepared in 1989. Buffalo states “[t]he planning process will use the Wyoming BLM Mitigation Guidelines to develop management options and alternatives” (Buffalo Field Office, High Plains District 2015*a*, p. 9). Both Little Snake’s and Buffalo’s decisions state that they “will strive to be compatible with the existing plans and policies of adjacent” agencies (Little Snake Field Office 2011, p. RMP-6), (Buffalo Field Office, High Plains District 2015*a*, p. 7). Colorado River Valley’s “will be compatible” with them (Colorado River Valley Field Office 2014, p. 1-13). That said, the Wyoming guidelines add that “[c]onsistency in this sense does not mean that identical requirements would be applied for all similar types of land use activities” (Casper Field Office, Wyoming State Office 2007, Appendix I, p. I-1).

D.4 A Standard as the Answer to Both Sides

This section provides examples supporting the assertion that a buffer taken from an established standard can be defended to public stakeholders (both pro-conservation and pro-development) by citing the standard.

Raptor nests: Buffalo, Wyoming. Buffalo’s alternatives included a 1.5-mile seasonal buffer around the nests of special status raptors in the alternative the plan labels “Resource Conservation” and a quarter-mile buffer in the one it labels “Resource Utilization” (Buffalo Field Office, High Plains District 2015*b*, p. 217). The comments split. “Some commenters felt the raptor spatial buffer distances and dates were too restrictive while others commented that the protections were not sufficient” (Buffalo Field Office, High Plains District 2015*b*, p. 30). The office did not defend a distance of its own. It revised the raptor actions “to fully conform to the distances and dates recommended by the USFWS Wyoming Ecological Services Office” (Buffalo Field Office, High Plains District 2015*b*, p. 30). In its formal response, the part that concerns buffers is one sentence long and serves both groups. Commenters who called the nest buffers “unnecessary and not adequately justified” and commenters who “proposed additional restrictions protecting eagles” received the same answer: “Dates and distances for buffers were based on recommendations from the appropriate state and federal agencies” (Buffalo Field Office, High Plains District 2015*b*, p. 2716).

Raptor nests: Grand Junction, Colorado. Three commenters pulled Grand Junction’s raptor buffers in three directions. An individual asked the office to go “beyond the 0.33 mile” around ferruginous hawk

nests. The county commissioners objected that the Fish and Wildlife Service “recommends only a 200 meter” buffer around eagle nests and that the office “has not justified the need for this additional buffer.” The state wildlife agency asked for no-surface-disturbance stipulations in place of seasonal limits. The office gave one answer to all three: “Buffers for raptor stipulations were developed to be consistent with the best available science, including CPW’s Recommended Buffer Zones and Seasonal Restrictions for Colorado Raptors (CPW 2008). No change has been made” (Grand Junction Field Office 2015, p. 6-45). The county had cited a standard of its own; the one published by the state’s wildlife agency prevailed.

Sensitive plants: Grand Junction, Colorado. Asked to defend a 200-meter buffer around sensitive plant communities, Grand Junction answered with outside references, none of them specific to its planning area. The justification “comes from Onshore Oil and Gas Order No. 1,” a national regulation. The distance “is consistent with NSO stipulations in adjacent BLM field offices.” It “is supported by” a 2009 set of recommended practices for rare plants and by draft guidance from the Fish and Wildlife Service (Grand Junction Field Office 2015, p. 6-280). NSO, the plans’ abbreviation in these quotations, is no surface disturbance: the regime that bars drilling equipment from the surface.

Historic sites: Buffalo, Wyoming. Commenters told Buffalo that its stipulations around cultural sites “were too restrictive.” The office answered that the three-mile distance rested on the state historic preservation office’s “experience with historic trails issues in other parts of Wyoming,” and concluded: “The distance is not arbitrary since it is consistent with other Wyoming BLM plans” (Buffalo Field Office, High Plains District 2015b, p. 2722).

Wildlife stipulations: Uncompahgre, Colorado. Uncompahgre states the practice as a rule and applies it in both directions. TLs are timing limitations, restrictions that apply for part of the year. Commenters asked the office to acknowledge that these “are not adequate to protect big game species from oil and gas development” (Uncompahgre Field Office 2019, p. R-220). The office answered: “The BLM defers to existing federal and state guidelines and policy, as well as agency expertise, such as CPW’s and USFWS’, to help craft effective land use stipulations.” Where the state agency’s version does not fit, the plan “ensures TLs are consistent with BLM statewide stipulations” (Uncompahgre Field Office 2019, p. R-223). A rural electric association asked for the opposite, that timing limitations be “the minimum biologically necessary” (Uncompahgre Field Office 2019, pp. R-217, R-221). It was told that they “were developed based on best available biological information in coordination with CPW” (Uncompahgre Field Office 2019, p. R-224).

D.5 Departures from a Standard

This section provides examples supporting the assertion that a buffer not taken from a standard has to be defended on its own. In each case a departure from a standard is what draws the demand for data, or what the office defends with local facts.

Lease stipulations: Uncompahgre, Colorado. Colorado’s wildlife agency told Uncompahgre that some stipulations in the preferred alternative were “not consistent with lease stipulations recommended to BLM by CPW.” The agency then set the terms for a departure: “If the BLM elects to not incorporate our recommendations in the preferred alternative, we recommend the analysis provide data, with clear rationale supporting that decision” (Uncompahgre Field Office 2019, pp. R-209–R-210).

Fish-bearing waters: Buffalo, Wyoming. Buffalo placed a quarter-mile restriction around fish-bearing waters, wider than the 500-foot standard. Commenters objected. The office’s defence names the standard and gives a physical reason for exceeding it: “the buffer from Onshore Order No. 1 is insufficient to protect the fish habitat from sedimentation since canyon rims are well over 500 feet distance from the water” (Buffalo Field Office, High Plains District 2015b, p. 2717).

Major rivers: Grand Junction, Colorado. Grand Junction proposed buffers around major rivers wider than the usual stream buffer. An industry group asked for the “scientific justification,” noted that “[h]istoric buffers have been limited to 300 to 500 feet,” and argued that “[w]ithout scientific evidence proving its need, this new restriction is unwarranted” (Grand Junction Field Office 2015, pp. 6-270–6-271). The office replied at length. It cited the statewide working group, a statewide risk assessment for drinking water, documented contamination incidents, and “the topography and river channel type in this planning area” (Grand Junction Field Office 2015, p. 6-271).

Raptor nests: Rock Springs, Wyoming. A county objected that a two-mile raptor buffer in Rock Springs’ draft “is more restrictive than United States Fish and Wildlife standards,” and that “[t]he BLM has not provided information to support a 2-mile buffer zone.” Other commenters wrote that the Service’s “recommended distances should be used.” The office’s replies on that page do not defend the distance. Each refers the commenter to the range of alternatives (Rock Springs Field Office, High Desert District 2024*a*, p. W-149). The approved plan sets raptor buffers of a quarter mile to one mile by species (Rock Springs Field Office, High Desert District 2024*b*, p. 2-32), and the decision record gives alignment with adjacent offices as a reason for the final stipulations (Rock Springs Field Office, High Desert District 2024*b*, pp. 1-4–1-5).

D.6 Standards That Are State-Specific and Do Not Bind

This section provides examples supporting the assertion that just because a standard exists does not mean it is followed. A standard, as the paper uses the word, is a published buffer an office can cite: a state agency’s guidance, another office’s choice, or the office’s own earlier plan. The examples in this section show offices declining standards.

Grand Junction, Colorado. A conservation group asked Grand Junction to adopt “the same buffer setback stipulations” that Little Snake had approved in 2011, a quarter mile around all perennial waters (Grand Junction Field Office 2015, p. 6-276). The office kept its own buffers and replied that its analysis “did not show that larger buffers were warranted” (Grand Junction Field Office 2015, p. 6-280). Grand Junction told commenters it “is not obligated to be perfectly consistent with CPW’s guidelines” (Grand Junction Field Office 2015, p. 6-348).

Little Snake, Colorado. A conservation group asked Little Snake to adopt raptor guidelines from the Fish and Wildlife Service’s Utah office. Little Snake answered that it “is not required to use guidelines prepared by another ecological field office” (Little Snake Field Office 2010, p. Q-247).

Uncompahgre, Colorado. Uncompahgre states “BLM state-specific guidance outside of Colorado are not applicable” (Uncompahgre Field Office 2019, p. R-68). When Colorado’s Governor objected that Uncompahgre had not adopted the state agency’s stipulations, the BLM replied that “such recommendations do not comprise an officially approved policy or program” under the consistency regulation (Bureau of Land Management, Colorado State Office 2020, pp. 2–3).

E Calibrating the Ecological Ranges

This section provides the evidence behind the effective ecological range h Section 3.3 defines.

E.1 The Range of Each Feature Family

This section lists the range of every feature family in the estimation sample.

Water features. The range is 120 m for riparian wetlands and for floodplains, whose mapped extent is itself the strip that shields the water, so the range covers only the ground that drains into it; 200 m for streams, rivers and waterbodies, and by analogy for drinking-water sources; 400 m for the named major rivers; and 805 m, half a mile, for groundwater. Section E.2 derives the 200 m.

Wildlife features. The range is 400 m for prairie dog towns; 1,609 m, one mile, for bald eagle roosts; and 2,414 m, a mile and a half, for nests. Section E.3 derives the nest range. For prairie dog towns and bald eagle roosts, whose sites are areas rather than points, the range prices only the ring outside the mapped site, since every acre inside it carries full weight whatever the range (footnote 33 of Section 3.3).

Soil features. The four soil families carry no compact geometry to measure distance from, so each is given two anchors in attribute space rather than metres: the level at or above which an avoided acre carries full conservation value, and the level below which it carries essentially none. The shortfall between them is the range h . For slope the anchors are a 50% and a 10% grade; for water erodibility a RUSLE factor of $K = 0.64$ and $K = 0.20$; for wind erodibility an erodibility index of 310 and 56 t/ac/yr; and for saline soils an electrical conductivity of 16 and 8 dS/m. Section E.4 derives the slope anchors.

E.2 Distance Features: Streams

This section derives the 200 m range of streams, rivers and waterbodies. An acre at the water's edge earns full weight, decaying to near zero by about 200 m from the edge. Two distances set this range, and the range takes the upper end of each.

The riparian strip. The first distance is the width of the strip along the stream that keeps its ecosystem healthy: the ground that filters runoff before it reaches the water, holds the bank in place, and supplies streamside habitat for wildlife. Buffer design guidelines written for this region put that strip at up to about 90 m (Johnson & Buffler 2008, p. 26).

The overland path. The second distance is how far upslope of that strip surface disturbance can still affect it. The federal rangeland erosion model puts the distance runoff travels before it soaks into the ground or drops the sediment it carries at no more than about 120 m (U.S. Department of Agriculture, Natural Resources Conservation Service 2021).⁵²

Why the upper end of each. What the weighting needs is the distance at which an avoided acre stops mattering *altogether*, so each input is how far its mechanism can plausibly reach, not how far it typically does. Past roughly 200 m, an acre neither belongs to the strip that shields the stream nor drains into it, so disturbance beyond this distance is given almost no weight.

E.3 Distance Features: Nests

This section derives the 2,414 m range of raptor nests, a mile and a half. The range is the outer edge of the federal guidance that grades restrictions near an active nest by distance. The Forest Service's conservation assessment for the ferruginous hawk recommends a maximum duration of disturbance by breeding stage and by distance from the nest, in three bands: within half a mile, half a mile to one mile, and one mile to a mile and a half. "Out-of-vehicle" disturbance, which the schedule defines to include "work-related activities (e.g., mineral development, agricultural operations)," is permitted for no time at all inside half a mile during the sensitive breeding stages, for under an hour in the second band, and for under a day in the third. The

⁵²The 120 m is the longest slope the model supports, not a distance sediment was observed to move; its standard case is 50 m (U.S. Department of Agriculture, Natural Resources Conservation Service 2021, 647-C.3). Sediment typically travels only a fraction of that distance—monitored plumes at gas well pads extend under 40 m beyond the disturbed area (Williams et al. 2008)—so 200 m is an outer bound rather than a tight estimate.

schedule states no restriction beyond a mile and a half (Collins & Reynolds 2005, Table 7, p. 44). A well pad, with its road and its production life, is permanent, and so exceeds every one of these caps. On this schedule a pad is incompatible with an active nest anywhere inside a mile and a half and unrestricted beyond it. So I set the ecological range at 1.5 miles.

The range is an upper bound. Every other distance the assessment attaches to human disturbance near a nest lies inside this one. It recommends posting wildlife alert signs at least 450 m from active nests (Collins & Reynolds 2005, p. 43), and placing artificial nest platforms at least one mile from any public road and at least a mile and a half from any occupied building (Collins & Reynolds 2005, p. 42). The schedule’s own bands end at a mile and a half, and no recommendation in the assessment reaches farther. The range is therefore the farthest distance at which the cited guidance treats disturbance as affecting a nesting pair, an upper bound rather than a typical response distance.

E.4 Threshold Features: Slopes

This section derives the two anchors of the slope range. In the distance case the maximum weight of one sits at the feature and decays with distance until disturbance affects the feature negligibly. In the threshold case it is less obvious where the maximum weight should land, so two points are calibrated: a *lower anchor*, below which disturbance does little lasting harm, and an *upper anchor*, above which almost any disturbance is harmful. For slope the anchors come from the two distinct processes by which disturbing a slope degrades it, surface erosion and mass wasting.

The lower anchor: surface erosion. On gentle slopes the concern from disturbance is surface erosion. Many factors set how fast a disturbed slope sheds soil—rainfall, soil erodibility, ground cover—and steepness adds to that loss at every grade, not only steep ones. What the federal erosion model marks at about 9% is a change in *how fast* steepness adds erosion: below that grade an extra degree of slope adds little, above it each degree adds soil loss more sharply (Renard et al. 1997, ch. 4). A grade near 10% therefore marks where slope begins to drive erosion sharply, the point at which avoiding disturbance begins to meaningfully benefit soil health.

The upper anchor: mass wasting. The upper anchor is where the more damaging process, mass wasting (landslide, slope failure), becomes markedly more likely. The Forest Service’s road-failure records show “a significant threshold at 50 percent,” with mass-wasting failures per mile almost eight times higher above that grade than below (U.S. Department of Agriculture, Forest Service 1999, App. 2).

F Development Forecasts

Section 3.3 builds the payoff attributes from per-acre revenue and disturbance densities, $m_{i,c}$ and $d_{i,c}$, defined over each resource’s development-potential categories c . This appendix details how those densities are recovered from the planning record.

Forecasts. Regulation requires each plan to include a *Reasonably Foreseeable Development* (RFD) report (Bureau of Land Management 1985): projections of commercial development under a reference scenario in which development is permitted everywhere legally allowed, reported at coarse spatial scales—typically by development-potential category or township. For oil and gas, the projections include the expected number of wells, success rates for exploratory (“wildcat”) and development wells, expected production per well, and associated surface disturbance. For each category, the RFD expresses expected development as a rate referenced to a township, a fixed areal unit (36 sq. mi.) much smaller than the category’s full extent—a category may be labeled “> 150 wells per township.” Converting this rate to a per-acre basis and multiplying by the expected disturbance and revenue per well gives each category its disturbance density $d_{i,c}$ (disturbed

acres per acre) and revenue density $m_{i,c}$ (revenue per acre) under the baseline scenario; both incorporate the probability of commercially viable extraction as the RFD reports it. Integrating a density over a footprint treats development as uniform within each category—the township-referenced rate applied evenly to every acre of the category—which is consistent with the coarse resolution at which RFDs report and imposes no further structure on where development falls within a category.

Regime shares. To evaluate restrictive regimes, expected outcomes are parameterized as proportional reductions against the baseline: for each regime r , $d_{i,c}(r) = a_{r,c} d_{i,c}$ and $m_{i,c}(r) = a_{r,c} m_{i,c}$, where $a_{r,c}$ is the share of the wells the baseline projects in category c that regime r allows. A single share scales both densities because both derive from the same projected well count: a restriction raises operating costs and so lowers the probability that a projected well is drilled, and expected disturbance and expected revenue fall with the well count together. By construction $a_{\text{Open},c} = 1$, with $a_{r,c}$ weakly decreasing in policy stringency,

$$1 = a_{\text{Open},c} \geq a_{\text{Discretionary},c} \geq a_{\text{Closed},c} = 0.$$

The record’s five stipulations collapse onto these three as in Section 3.2: timing limitations and controlled surface use share the discretionary factor, and no surface disturbance and Closed both permit no wells on the footprint. Where a plan publishes no reduction for its discretionary stipulations, the share is borrowed from the plans that do, interpolated on the resource’s own well density. Where more than one fluid mineral resource is viable at a feature, the feature-level quantities sum the viable resources’ expectations—each read off that resource’s own potential map under the same pair $\mathbf{y}$ —so $M_{i,f}$ and $K_{i,f}$ price the one decision against the combined development it restricts.

Table A1: Willingness to Pay for Conservation by Field Office and Feature Group: Conditional Logit Specification

Field office	Willingness to pay (\$/wtd. acre)			Policy fixed effects (logits)		Scale (log σ)		Summary statistics		
	Soil	Water	Wildlife	Leaves all open	Any closed label	Soil+Water	Wildlife	McFadden R^2	log L	Features
Pooled (all offices)	\$76,144***/\$59,511*** (\$7,071)/(\$6,531)	\$34,488***/\$30,777*** (\$8,065)/(\$7,012)	\$44,346***/\$31,815*** (\$1,775)/(\$1,724)	0.927***/0.930*** (0.0254)/(0.0254)	-0.553***/-0.549*** (0.0213)/(0.0214)	5.68/5.69 (0.0806)/(0.0827)	3.74/3.78 (0.0492)/(0.0525)	0.058/0.057	-35,310/-35,352	12,509
Rock Springs, WY	\$57,357***/\$48,716*** (\$16,903)/(\$10,347)	\$47,495***/\$42,168*** (\$6,543)/(\$5,065)	\$36,093***/\$30,572*** (\$509)/(\$286)	0.0535/0.0717 (0.0847)/(0.0845)	-11.1***/-9.25*** (2.19)/(1.43)	3.36/3.29 (0.165)/(0.164)	-0.673/-0.901 (0.0667)/(0.0686)	0.334/0.336	-2,659/-2,654	1,222
Casper, WY	\$93,167***/\$59,380*** (\$7,742)/(\$4,658)	\$4,915/\$6,431* (\$4,198)/(\$3,400)	\$22,058***/\$16,393*** (\$709)/(\$592)	-0.109/-0.136* (0.0735)/(0.0739)	-2.11***/-2.09*** (0.0806)/(0.0802)	3.72/3.66 (0.129)/(0.128)	1.15/1.14 (0.0595)/(0.0592)	0.182/0.180	-5,099/-5,108	2,036
Buffalo, WY	\$36,117***/\$31,417*** (\$3,285)/(\$2,315)	\$5,957/\$10,582** (\$6,847)/(\$5,316)	\$59,229***/\$43,653*** (\$1,386)/(\$929)	-0.514***/-0.530*** (0.0765)/(0.0763)	-1.84***/-1.83*** (0.0562)/(0.0562)	3.75/3.63 (0.128)/(0.129)	2.99/2.86 (0.0459)/(0.0456)	0.142/0.144	-9,411/-9,389	3,300
Lander, WY	\$261,324***/\$172,990*** (\$36,086)/(\$22,341)	\$62,930***/\$65,723*** (\$21,440)/(\$18,086)	\$122,483***/\$99,128*** (\$29,500)/(\$15,151)	1.26***/1.26*** (0.117)/(0.117)	-0.436***/-0.433*** (0.103)/(0.103)	5.01/4.95 (0.181)/(0.182)	6.16/5.80 (0.411)/(0.365)	0.111/0.112	-1,562/-1,560	577
Grand Jct., CO	\$189,666***/\$126,141*** (\$30,382)/(\$21,147)	\$556,687***/\$426,777*** (\$96,757)/(\$78,270)	\$20,170*/\$15,860 (\$11,556)/(\$10,587)	0.214/0.233 (0.258)/(0.258)	0.875***/0.923*** (0.151)/(0.150)	5.19/5.38 (0.194)/(0.205)	-0.0957/-0.0765 (0.457)/(0.489)	0.136/0.126	-628/-635	267
Uncompahgre, CO	\$23,647***/\$24,053*** (\$2,806)/(\$2,157)	\$106,720***/\$78,398*** (\$12,648)/(\$7,278)	\$40,901***/\$30,768*** (\$1,640)/(\$1,675)	1.65***/1.63*** (0.143)/(0.144)	0.668***/0.713*** (0.119)/(0.118)	2.98/2.76 (0.203)/(0.191)	-0.777/-0.275 (0.163)/(0.183)	0.149/0.137	-1,158/-1,173	480
Royal Gorge, CO	\$447,008**/\$365,537*** (\$177,985)/(\$85,039)	\$350,542***/\$347,256*** (\$24,823)/(\$23,568)	\$1,097,580***/\$788,737*** (\$30,252)/(\$17,400)	2.08***/2.08*** (0.0652)/(0.0652)	1.16***/1.18*** (0.0578)/(0.0577)	4.96/4.77 (0.178)/(0.168)	0.643/1.26 (0.143)/(0.169)	0.125/0.119	-5,242/-5,276	2,350
Colo. River Valley, CO	\$75,592***/\$72,405*** (\$4,591)/(\$3,053)	\$65,422***/\$66,153*** (\$6,870)/(\$4,921)	\$121,574***/\$88,481*** (\$1,251)/(\$759)	1.98***/1.99*** (0.169)/(0.170)	1.74***/1.79*** (0.141)/(0.141)	4.46/4.26 (0.250)/(0.254)	-0.940/-0.612 (0.133)/(0.134)	0.261/0.239	-1,077/-1,109	492
Little Snake, CO	\$8,116/\$7,262 (\$6,237)/(\$4,588)	\$18,308**/\$14,433*** (\$7,268)/(\$5,359)	\$22,063***/\$16,435*** (\$150)/(\$82)	2.85***/2.81*** (0.158)/(0.156)	0.745***/1.19*** (0.165)/(0.155)	4.50/4.24 (0.226)/(0.207)	-2.21/-2.07 (0.0816)/(0.0780)	0.542/0.470	-925/-1,071	650
Four Rivers, ID	\$43,674***/\$45,340*** (\$11,394)/(\$8,577)	\$243,355***/\$180,912*** (\$22,300)/(\$15,572)	\$5,507/\$4,145 (\$4,017)/(\$2,954)	1.41***/1.41*** (0.131)/(0.131)	0.395***/0.408*** (0.112)/(0.112)	0.795/0.744 (0.164)/(0.165)	-5.37/-5.41 (0.467)/(0.502)	0.116/0.114	-1,204/-1,207	499
Rio Puerco, NM	\$15,852*/\$12,600* (\$9,620)/(\$7,408)	\$21,565**/\$21,444** (\$9,111)/(\$8,593)	<i>no data</i>	2.19***/2.19*** (0.0946)/(0.0946)	-1.10***/-1.10*** (0.137)/(0.137)	2.32/2.28 (0.430)/(0.429)	<i>no data</i>	0.290/0.290	-1,162/-1,162	636

Notes: Table reports the estimates from a conditional logit specification—the main specification setting $\kappa = 0$ and dropping the open type. Each office row is fit on that office alone; the pooled row fits the same specification once on all eleven offices together, a single restricted fit rather than an average of the eleven. Each cell gives the estimate under relevance curvature $p = 1.5$ before the slash and $p = 2$ after it, so that an acre at distance l from the feature carries relevance $0.01(l/h)^{1.5}$ or $0.01(l/h)^2$; the two differ only in how sharply relevance falls away with distance from the feature or, for soil, in attribute space. Willingness to pay is the revenue, in 2014 dollars, an office forgoes to protect an acre at the feature from development, for each natural-capital class. The policy fixed effects are the office’s residual tendencies to leave every resource open and to close a feature outright, each measured against the discretionary stipulations short of closure, common to all natural-capital classes and measured in log-odds rather than dollars. The scale, reported in logs for soil and water together and for wildlife, is the spread of the idiosyncratic differences among a feature’s alternatives. The summary statistics are McFadden’s R^2 , the log-likelihood and the number of features in the fit. *no data* marks a class×office group that I do not currently have data for. Standard errors in parentheses; the conditional logit is written in willingness-to-pay space, so the willingness-to-pay standard errors are ordinary rather than delta-method. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.


Table A2: Correlation Between Revenue Forgone and Disturbance Avoided Across a Feature’s Alternatives

	Features	Disturbance Avoided with All Acres Weighted Equally						Disturbance Avoided with Relevance-Weighted Acres					
		All alternatives			Within regime			All alternatives			Within regime		
		Median	p10	p5	Median	p10	p5	Median	p10	p5	Median	p10	p5
Rock Springs, WY	1,140	0.999	0.986	0.968	1.000	0.989	0.974	0.887	0.734	0.652	0.859	0.670	0.556
Casper, WY	1,973	1.000	0.998	0.995	1.000	0.998	0.996	0.853	0.701	0.652	0.841	0.639	0.578
Buffalo, WY	3,190	1.000	0.998	0.993	1.000	0.998	0.994	0.875	0.751	0.722	0.849	0.645	0.613
Lander, WY	487	0.999	0.990	0.967	1.000	0.990	0.971	0.842	0.638	0.599	0.815	0.481	0.463
Grand Jct., CO	222	1.000	0.985	0.970	1.000	0.994	0.987	0.888	0.669	0.602	0.854	0.519	0.460
Uncompahgre, CO	418	1.000	0.995	0.988	1.000	0.996	0.989	0.864	0.720	0.694	0.841	0.657	0.613
Royal Gorge, CO	1,856	1.000	0.989	0.967	1.000	0.991	0.976	0.829	0.537	0.507	0.797	0.519	0.471
Colo. River Valley, CO	412	1.000	0.996	0.965	1.000	0.996	0.985	0.876	0.677	0.622	0.858	0.549	0.477
Little Snake, CO	602	1.000	0.999	0.998	1.000	1.000	0.998	0.886	0.730	0.701	0.863	0.626	0.571
Four Rivers, ID	391	0.996	0.967	0.944	1.000	0.985	0.955	0.870	0.642	0.603	0.832	0.504	0.458
Rio Puerco, NM	536	1.000	0.994	0.980	1.000	0.995	0.982	0.800	0.603	0.558	0.841	0.511	0.471
Pooled (all offices)	11,227	1.000	0.994	0.983	1.000	0.996	0.986	0.868	0.669	0.597	0.847	0.591	0.515

Notes: Entries are the correlation, within each feature’s choice set, of revenue forgone M with disturbance avoided K . Columns report the median across features and the 10th and 5th percentiles, the less collinear tail. *All alternatives* uses every alternative in the feature’s choice set—the variation the choice model is estimated from. *Within regime* uses only alternatives that share the same regulatory regime—for example, all alternatives that close the footprint—so they differ in the extent of the footprint alone. *Disturbance avoided with relevance-weighted acres* is K as it enters estimation. *Disturbance avoided with all acres weighted equally* is computed the same way, with every acre’s relevance weight set to one. With all acres weighted equally, M and K are nearly proportional by construction: a cell’s expected disturbance and its expected revenue are both set by the cell’s development-potential level, so over any footprint both are sums of the same acres by level, and they diverge only as the mix of levels changes with extent. The relevance weight enters disturbance avoided and never revenue, and is what separates the two. Features with fewer than 3 distinct extents, or with no revenue at stake, are omitted (1,282 of 12,509).

Figure A1: Example Benefit-Cost Analysis Inputs from Rock Springs, Wyoming

(a) Oil and Gas Development Potential



(b) Projected Annual Resource Outputs

--	Alt. A	Alt. B	Alt. C	Alt. D
Livestock Grazing (AUMs – Billed Use) ¹	147,631	147,631	147,631	147,631
Livestock Grazing (AUMs – Total Authorized Use) ^{1,4}	303,238	297,066	160,387	303,238
Conventional Oil and Gas Wells Drilled ¹	232.4	61.5	238.8	230.2
CBNG Wells Drilled ¹	6.2	3.0	7.2	6.7
Crude Oil (bbbls – 2016) ²	2,047,125	542,141	2,104,317	2,026,774
Natural Gas (mcf – 2016) ²	82,824,048	21,947,512	85,144,702	82,005,381
Coal and Soda Ash (short tons) ^{1,3}	4,107,267	4,107,267	4,107,267	4,107,267
Recreation (visits – Low Visitation) ¹	426,439	426,439	426,439	426,439
Recreation (visits – High Visitation) ¹	847,318	847,318	847,318	847,318

(c) Big game habitat by Regulatory Regime

	Alternative A	Alternative B	Alternative C	Alternative D	Proposed RMP
Oil and Gas—Closed					
Elk	561,654	2,713,522	225,782	493,425	220,204
Elk Parturition	92,351	137,100	71	23,422	108,373
Mule Deer	561,656	2,713,530	225,782	493,430	182,445
Mule Deer Parturition	54,722	63,113	4,011	4,011	55,214
Moose	124,815	773,640	24,237	134,474	7,972
Pronghorn	561,656	2,716,103	225,782	496,002	159,583
Oil and Gas—NSO					
Elk	518,268	809,743	15,542	194,648	20,094
Elk Parturition	31,765	6,246	3,094	8,705	4,653
Mule Deer	518,270	809,800	15,542	194,648	11,517
Mule Deer Parturition	3,314	1,568	6,128	27,201	4,395
Moose	176,191	234,198	38,623	126,673	144
Pronghorn	518,270	815,956	15,542	194,648	52,517
Oil and Gas—CSU					
Elk	1,957,092	126,442	1,924,802	2,410,222	70,860
Elk Parturition	17,487	0	47,155	106,998	23,542
Mule Deer	1,957,104	126,426	1,924,917	2,410,233	229,419
Mule Deer Parturition	6,450	0	25,783	33,366	4,473
Moose	502,660	26,580	627,145	581,884	11,190
Pronghorn	1,978,008	127,383	1,929,014	2,418,158	254,216
Oil and Gas—TLS					
Elk	2,600,380	970,225	2,917,087	2,279,364	79,436
Elk Parturition	50,995	6,246	143,275	119,924	30,249
Mule Deer	2,600,433	970,257	2,917,141	2,279,435	184,835
Mule Deer Parturition	9,959	1,568	60,670	60,670	5,024
Moose	792,831	240,618	896,598	801,754	6,018
Pronghorn	2,602,246	982,994	2,918,958	2,731,596	140,448

(d) Projected Surface Disturbance (Alt. C)

Disturbance Associated With All New Drilled Wells and Existing Active Wells (Short-Term Disturbance)					
Wells			Acres of Surface Disturbance		
Well Type	All Lands	BLM Managed	Well Pad + Access Roads + Pipelines	All Lands	BLM Managed
New Conventional Oil and Gas	5,902	4,804	7.1 successful and 5.6 for drilled and abandoned	40,575	33,024
New Coalbed Natural Gas	220	141	7.1 successful and 5.6 for drilled and abandoned	1,545	990
Existing Conventional Oil and Gas	2,259	1,536	2.64	5,964	4,055
Existing Coalbed Natural Gas	52	28	2.64	137	74
Total Wells	8,433	6,509	Total Disturbance	48,221	38,144

Notes: Each panel reproduces a planning-document table or map projecting development and conservation under different zoning configurations. Panel (a) is an example development-potential map for oil and gas, partitioning the planning area into ordered categories from lower to higher expected development in a zoning configuration where development is allowed everywhere the field office can legally permit it; see Appendix F for more details. Panel (b) shows the projected annual resource outputs tabulated by intersecting each considered configuration with development-potential maps for each resource (see Panel (a) for an example) (Rock Springs Field Office, High Desert District 2024a). Panel (c) shows the big game habitat in closed and stipulated lands, by species, highlighting the primary means conservation is factored into the benefit-cost analysis (acreage) (Rock Springs Field Office, High Desert District 2024a). Panel (d) shows the surface disturbance projected under one configuration over the plan period, obtained by multiplying the wells that configuration's zoning is expected to permit by an average per-well footprint of well pad, access roads, and pipelines. The report projects those well counts by working out how many well locations each configuration's restrictions remove from a baseline in which drilling proceeds wherever law allows, and tabulates the resulting disturbance separately for the baseline and for each alternative (Stilwell 2013).

Figure A2: Example Benefit-Cost Analysis Inputs from Four Rivers, Idaho: Permitted Disturbance by Vegetation Community


Table 4.4.2 Estimated acres of permitted disturbance for each alternative by vegetation community

Plant Community	Acres (total)	Acres of Human Caused Disturbance Not Including Rehabilitation or Restoration ²			
		Alt. A	Alt. B	Alt. C	Alt. D
Evergreen Forest	50,830	940	600	1,500	1,000
Aspen & Mountain Shrub	47,150	1,300	1,000	1,300	1,200
Shrub-Steppe & Perennial Grassland	573,110	38,430	32,000	42,000	39,000
Exotic Annual Grassland	162,750	13,700	16,000	15,000	15,000
Riparian & Wetland	38,140	3,010	2,000	3,000	2,500

²Surface disturbance from permitted activities (new ROW, new rock & mineral extraction, new roads & trails, commercial logging/firewood extraction, etc.).

Notes: Reproduces Table 4.4.2 of the Four Rivers Field Office Draft RMP and Draft EIS (Four Rivers Field Office, Boise District 2019). Each row is a vegetation community. The first column gives each community's total acreage in the planning area, and the remaining columns give the acres where the planning documents project human-caused surface disturbance under each zoning alternative (Alt. A through Alt. D).


Figure A3: Willingness to Pay for Conservation, by Field Office and Feature Group: Conditional Logit Specification



Willingness to pay to avoid disturbing the feature's most ecologically valuable acreage (\$)

Notes: Figure plots each field office's estimated willingness to pay for conservation, in dollars of forgone revenue, from a conditional logit specification—the main specification setting $\kappa = 0$ and dropping the open type. Each bar is the estimate for one office and one natural-capital class under relevance curvature $p = 2$, so that an acre at distance l from the feature carries relevance $0.01^{(l/h)^2}$; the estimates under $p = 1.5$ are in Table A1. The whiskers are 90% confidence intervals from the Hessian standard errors. Each office is fit on its own; the *Pooled* bars fit the same specification once on all eleven offices together. The policy fixed effects and the scales are not plotted; Table A1 reports them. The horizontal axis is signed-log, matching Figure 4. “no data” marks an office×class group with no data.


Figure A4: Willingness to Pay for Conservation, by Field Office and Feature Group: Mixed Logit at Relevance Curvature $p = 1.5$



Willingness to pay for conservation (2014 dollars)

Notes: Figure plots each field office's estimated willingness to pay for conservation, in dollars of forgone revenue, from Section 2's model and Section 4.2's specification. Each bar is the estimate for one office and one natural-capital class under relevance curvature $p = 1.5$, so that an acre at distance l from the feature carries relevance $0.01^{(l/h)^{1.5}}$; the estimates under $p = 2$, the paper's specification, are Figure 4. The whiskers are 90% confidence intervals from the Hessian standard errors. The fit reaches a log-likelihood of $-23,530.63$ on the same 12,509 choice situations. The horizontal axis is signed-log, matching Figure A3. "no data" marks an office $\times$ class group with no data.

Figure A5: Willingness to Pay for Conservation and the 2016 Trump Vote Share: Precincts Weighted by BLM Acres




Notes: Figure repeats Figure 5 with a different measure of the vote share. Each marker is one office, and its shape denotes the office's state. The panels include the twenty-eight estimates that are statistically distinguishable from zero. Here each precinct's two-party Trump share is weighted by the acres of BLM-managed surface the precinct holds inside the planning area, so that precincts without BLM land carry no weight. The vote share data is from Voting and Election Science Team (2018). Lines are least-squares fits of log willingness to pay on the vote share for each resource type: the dark line leaves out Rio Puerco, New Mexico (hollow marker), and the light line includes it. Rio Puerco is set apart because it is the only New Mexico office, and New Mexico's voters lean Democratic while its state government depends heavily on revenue from federal onshore mineral leasing (New Mexico Legislative Finance Committee 2023, Volcovici & Renshaw 2021). Rio Puerco has no wildlife estimate, so the wildlife panel has one line. The vertical axis is on a log scale, and the axis limits are those of Figure 5. The weighting moves ten of the eleven offices toward Trump, by 8 to 22 percentage points; Rio Puerco moves 6 points the other way.

Figure A6: The Fitted Discrete Distribution of Corrections for the Standard Proxy




Notes: Figure plots the estimated support of the correction η and each field office's mixing shares over it. The standard is measured with a proxy, the median relevance the other field offices chose for features of the same family; the correction allows the standard an office anchors on to sit short of that proxy or past it. The column headed *Proxy* is a correction of zero. Left of it the correction is positive and the office's standard stops short of the proxy, a narrower buffer; right of it the correction is negative and the standard reaches past the proxy. Bubble area is the office's mixing share for that correction; shares sum to one across a row, and shares below 0.1% are not drawn. Under each correction is the buffer the corrected standard implies for one example family, raptor nests: how far from the nest an office anchored on that standard protects, given the nests' proxy. A correction marked *past their ecologically relevant range* puts the standard beyond where a nest's relevance has fallen to nothing. Other families translate the same way, each relative to its own proxy. The left-hand column is the open type, which leaves a feature open regardless of the payoff; it is not a value of the correction and is placed at the left end of the axis. The bottom row averages the eleven offices with equal weight. The corrections describe the sample, not any one office: the effective number of corrections an office carries averages 1.4, and Buffalo, WY and Four Rivers, ID each carry a single correction. The open-type share is below 1% at 8 of the eleven offices; Colo. River Valley, CO is the only office carrying much, at 8%. Standard errors on the fitted corrections run from 0.0012 to 0.0080; they describe how sharply the fitted objective pins down each correction locally and are not confidence intervals for the location of a true latent value.

Figure A7: How Often the Predicted Buffer Yields Less Conservation Value Net of Revenue Forgone than an Adjacent Buffer




Notes: Figure plots the share of each field office’s decisions in which a buffer adjacent to the *predicted* buffer has a higher net benefit—its conservation value net of the revenue given up. Figure includes decisions for which the fitted model predicts protection: it assigns the protective alternatives together a higher probability than leaving the feature open. The *predicted* buffer is the protective alternative with the highest predicted probability. Its adjacent buffers are the alternatives under the same regime one step narrower and one step wider.

Figure A8: Predicted Closing With and Without the Closure Fixed Effect




Notes: Figure plots the share of protective decisions that close the feature rather than restrict it under a stipulation by field office. The black tick is the realized share. The filled gray marker is the full fitted model's predicted probability of closing conditional on being protected. The hollow marker predicts the probability with the fitted model but with each office's closure fixed effect $\gamma_i(\text{closed})$ set to zero. At the right is the office's closure fixed effect in log-odds with its standard error (Table 1).

Figure A9: Revising the Standard under Heterogeneous and Uniform Adoption




Notes: Figure plots changes in baby birds (fledglings) protected and revenue forgone, relative to the baseline, under heterogeneous and uniform adoption of the revised conservation standard. Dark orange bars show the standard revision under *heterogeneous adoption*, in which all offices respond to a common one-mile recommendation but retain office- and nest-specific unobserved heterogeneity in the conservation standards they ultimately follow. Light orange bars show the revision under *uniform adoption*, in which all offices respond to the same one-mile conservation standard without retaining this unobserved heterogeneity. The top panel reports aggregate changes, while the bottom panel breaks them out by office. The first column shows changes in fledglings protected, the second changes in revenue forgone, and the third, appearing only in the top panel, revenue forgone per additional fledgling. The “Other offices” category includes Uncompahgre, Colorado River Valley, Royal Gorge, and Grand Junction in Colorado, and Four Rivers in Idaho.

Figure A10: Protecting the Same Number of Birds with a Revised Standard (Uniform Adoption) or Price-Like Instruments



Notes: Figure replicates Figure 6 for the uniform adoption implementation of the standard revision counterfactual. It plots changes in baby birds (fledglings) protected and revenue forgone under the standard revision and price-like instrument counterfactuals, relative to the baseline. Light orange bars show the standard revision under *uniform adoption*, in which all offices respond to the same one-mile conservation standard, without retaining office- and nest-specific unobserved heterogeneity in the standards they follow. Light and dark blue bars show two implementations of the price-like instrument, both calibrated to protect the same aggregate number of additional fledglings as the standard revision. The *minimum value* implementation (light blue) raises each office's value of avoiding wildlife disturbance to at least \$284,000 per relevance-weighted acre, while the *common value* implementation (dark blue) sets every office's value to \$295,000. The top panel reports aggregate changes, while the bottom panel breaks them out by office. The first column shows changes in fledglings protected, the second changes in revenue forgone, and the third, appearing only in the top panel, revenue forgone per additional fledgling. The "Other offices" category includes Uncompahgre, Colorado River Valley, Royal Gorge, and Grand Junction in Colorado, and Four Rivers in Idaho.

Figure A11: Protecting Birds under a Common Value of \$120,000 per Fledgling



Notes: Figure plots each office's fitted value of a fledgling and the change in baby birds (fledglings) protected, relative to the baseline, when every office's value is set to \$120,000 per fledgling, the Fish and Wildlife Service's cost of offsetting one golden eagle. The left panel shows the value per fledgling each office already forgoes, on a log scale: the office's fitted value per relevance-weighted acre of avoided wildlife disturbance, times the disturbed acres per well in its plan (the median over its nest decisions), divided by two, since a well beside a nest costs a brood of two. The dashed line marks \$120,000. Hollow markers show the three offices whose wildlife value is not statistically distinguishable from zero. The right panel shows fledglings protected a season at the baseline (grey bars) and with every office set to \$120,000 per fledgling (colored bars), with the legend distinguishing offices that protect fewer fledglings from those that protect more. The counterfactual prices each nest decision at its own acres per well rather than the office's median, so an office near the line can move in either direction: Casper's median value is below \$120,000, yet it protects more. Nest decisions that expect no well have no price and are left at the baseline.