Borders, Spillovers, and the Watershed Lattice

Robert J. Carroll[*] David M. Konisky[†]
Christopher Reenock[‡]

March 2026

Preliminary draft. Comments welcome.

Abstract

We study the problem of drawing jurisdictional maps over a directed spillover network. Locations are nodes of a forest; $v \to w$ if a pollutant released at v reaches w downstream. A map is a partition of locations into tiles, one regulator per tile. We call a tile a *watershed tile* if it has a unique downstream outlet, and prove that the set of watershed tilings forms a Boolean algebra $\mathcal{W}(V, \to) \cong B_{|V|-k}$: drawing a map is equivalent to making $|V| - k$ independent binary boundary decisions. On any tiling the regulators play a local public goods game; we show the equilibrium aggregate composite is invariant across tilings (a neutrality result) and that equilibrium welfare is monotone on the Boolean lattice with direction determined by the sign of the spillover parameter α. We derive an anti-tone Galois connection between the tiling lattice and the space of intergovernmental grant schedules, establishing that map redesign and grant design are formally dual instruments. For any politically-drawn map $\mathcal{P}$, its watershed envelope $\mathcal{P}^{ws}$ weakly dominates it in welfare whenever $|\alpha|$ exceeds the bliss-point surplus ratio $\alpha^* = (b(\bar{\gamma}) - \bar{\gamma})/\bar{\gamma}$—where b is the benefit function and $\bar{\gamma}$ the equilibrium composite—giving a precise, computable measure of the cost of bad borders.

---

[*] Department of Political Science, University of Illinois at Urbana-Champaign. `rjc@illinois.edu`

[†] Paul H. O'Neill School of Public and Environmental Affairs, Indiana University. `dkonisky@iu.edu`

[‡] Department of Political Science, Florida State University. `creenock@fsu.edu`

1 Introduction

Pollution does not respect state lines, but regulators must. When a farmer in Pennsylvania applies nitrogen fertilizer, the runoff travels downstream through the Susquehanna River system and eventually degrades water quality in the Chesapeake Bay—a harm that falls on Maryland. Maryland regulators cannot compel Pennsylvania farmers to reduce applications; Pennsylvania regulators face no incentive to internalize harms that fall downstream. The externality is not a market failure in the usual sense: there are regulators on both sides, drawing their authority from the same federal system. The failure is jurisdictional: quite simply, the borders are in the wrong place.

This observation is as old as fiscal federalism. Oates (1972; 1999) argued that spillovers between jurisdictions are the canonical case for centralization, and that "rivers are the worst sort of border" precisely because water systems carry these spillovers directionally and irreversibly. But the argument is a heuristic. It does not say *how much* welfare is lost by drawing a river as a border, under what conditions the loss is large enough to justify redrawing the map, or whether redesigning the grants instead would achieve the same result. This paper provides all three.

We study the map design problem formally. A *spillover domain* $(V, \to)$ is a forest: locations are nodes, and $v \to w$ whenever a contaminant released at v reaches w downstream. A *jurisdictional map* is a partition of V into tiles. Each tile's regulator chooses how much to invest in pollution control; the investment generates a composite quality level that spills proportionally downstream. The provision game on any given tiling is the smooth-payoff version of the local public goods game studied in our companion paper (Carroll et al., 2025), which characterizes equilibrium on a fixed forest domain. Here we add the design dimension: a planner chooses the tiling before the game is played. Which tiling should she choose, and how does the answer depend on the magnitude and sign of the spillover parameter α?

Main results. The key tool is Birkhoff's representation theorem applied to the poset of admissible maps. The central observation it delivers: the set of *watershed tilings*—partitions in which each tile has a unique downstream outlet—forms a Boolean algebra $\mathcal{W}(V, \to) \cong B_{|V|-k}$ (Theorem 2.8). Drawing a jurisdictional map is equivalent to independently flipping $|V| - k$ binary switches, one per non-root location. This separability drives everything that follows.

Four results build on it. First, the equilibrium aggregate composite $\bar{\gamma}$ is invariant across all watershed tilings (Lemma 3.5): different maps redistribute the cost of provision across jurisdictions without changing the total. This is the directed-forest analogue of the Bergstrom-Blume-Varian (1986) neutrality theorem. Second, equilibrium welfare is monotone on the Boolean algebra with direction determined by the sign of α (Propositions 3.7 and 3.15): fragmentation is optimal when spillovers are positive ($\alpha > 0$), and consolidation is optimal when negative spillovers are strong enough ($|\alpha| > \alpha^*$). Third, there is an anti-tone Galois connection between the tiling lattice and the space of intergovernmental grant schedules (Proposition 3.11): for any map $\mathcal{P}$, the minimal grant schedule that compensates for bad borders is a closed-form function of the gap between $\mathcal{P}$ and its watershed comparator $\mathcal{P}^{ws}$. Map redesign and grant design are formally equivalent instruments, faces of the same duality. Fourth, watershed tilings dominate arbitrary political maps whenever $|\alpha| > \alpha^*$ (Proposition 3.15): restricting the design problem to $\mathcal{W}(V, \to)$ is without loss.

All four results are unified by a single structural constant. Writing $b(\cdot)$ for the benefit function each regulator derives from composite quality,

$$\alpha^* = \frac{b(\bar{\gamma}) - \bar{\gamma}}{\bar{\gamma}},$$

the *bliss-point surplus ratio*: the welfare surplus above provision cost at the optimal quality level, normalized by that level. It is the monotonicity threshold, the per-boundary break-even grant, and the watershed dominance condition. Whether any specific jurisdictional border is in the "Oates regime" reduces to the empirical question of whether the spillover transfer rate across that border exceeds α^*.

Policy implications. Three actionable conclusions follow directly from the structure of the model.

First: align borders with catchments, but not for the usual reason. The conventional case for watershed-aligned jurisdictions rests on the intuition that rivers "carry" the externality. The present results sharpen this: watershed alignment is not merely helpful but necessary for the design problem to be well-posed. Only watershed tilings form a Boolean algebra; only on a Boolean algebra do map design and grant design become dual instruments. Misaligned borders are not just suboptimal—they destroy the separability that makes the design problem tractable in the first place.

Second: the sign of the spillover determines the direction, but not by itself. A common presumption is that negative spillovers (pollution, runoff) call for merger and positive spillovers (knowledge, infrastructure) call for fragmentation. The results confirm the direction but add a condition: merger dominates only when $|\alpha| > \alpha^*$. A weak spillover, or a jurisdiction with high bliss-point surplus, can be in the positive-spillover regime even with nominally negative α. The practical recommendation is to measure α^* before recommending consolidation—it is computable from benefit-function data and equilibrium provision levels without any structural estimation.

Third: grants and borders are substitutes, at a price. When political constraints prevent redrawing the map, the Galois connection gives the minimum grant schedule that replicates the watershed outcome. The total cost $\sum_v \mu_v^* = \sum_v (|\alpha| - \alpha^*)$ per misaligned boundary is the fiscal price of the political constraint. This provides a direct way to evaluate programs like the Chesapeake Bay nutrient trading scheme: the program's budget should equal or exceed this threshold, and any shortfall identifies which boundaries are being left uncompensated.

Contributions. The mathematical tools used here—Birkhoff's representation theorem, Galois connections, closure operators—are classical. The contributions lie in their application and in the synthesis they enable.

Partition design has not been studied. There is a large and mature literature on centralized network design: a planner chooses which links to form in order to maximize equilibrium welfare (Belhaj et al., 2016; Baetz, 2015; Hiller, 2017; Li, 2023; Sun et al., 2026). The usual approaches in this literature take the node set as given and ask which edges to add. None of them asks how to partition a fixed network into jurisdictions—a question that is at least as natural, and arguably more tractable, for the broad class of applications where the underlying network (the river system, the power grid, the road network) is given by geography or infrastructure and the planner controls only administrative boundaries. The present paper opens this line of inquiry.

Tiling-invariance of the aggregate composite. Lemma 3.5 says that $\bar{\gamma}$—the equilibrium composite public good—is the same under every watershed tiling of the same domain. Different maps redistribute who bears the cost of provision but leave total quality unchanged. This is structurally analogous to BBV neutrality but operates through a different mechanism (forest best-reply dynamics, not income effects) and has a different design implication: welfare comparisons across tilings reduce entirely to cost comparisons.

The grant-tiling Galois connection. The fiscal federalism literature has long discussed map redesign and intergovernmental grants as alternative instruments for addressing spillovers, without establishing a formal relationship between them. Proposition 3.11 shows they are not merely alternatives but exact duals: the tiling lattice and the grant space are connected by an antitone Galois connection indexed by the same Boolean algebra. The duality is not an approximation or a special case; it holds exactly on forest domains and breaks on non-forest domains.

The threshold α^* and its unification. The same formula $(b(\bar{\gamma}) - \bar{\gamma})/\bar{\gamma}$ appears independently as the monotonicity threshold, the per-boundary break-even grant, and the watershed dominance condition. That α^* is network-independent is the content of Lemma 3.5: because every tile's equilibrium composite equals $\bar{\gamma}$ regardless of topology, the benefit term in every welfare comparison is always $b(\bar{\gamma})$ rather than some graph-dependent quantity. What does not follow from any prior result is that the *same* constant governs all three results simultaneously.

Related literature. The network design literature studies a planner who chooses which links to form among a fixed node set. Sun et al. (2026) organize this literature along two axes—centralized versus decentralized control, and static versus dynamic formation. The decentralized strand (Jackson & Wolinsky, 1996; Bala & Goyal, 2000) studies link formation as a non-cooperative game among agents; the centralized strand (Belhaj et al., 2016; Baetz, 2015; Hiller, 2017; Li, 2023; Sun et al., 2026) has a social planner choose which links to add to maximize equilibrium welfare. This paper is orthogonal to all of them: the network is fixed and the design variable is the partition of nodes into jurisdictions, not the set of edges.

The game on each watershed tiling is played on a forest domain and inherits the structure of (Carroll et al., 2025). The coarsest tiling $\hat{0}$ is a single tile and recovers the Bergstrom-Blume-Varian (1986) pure public goods game; the neutrality result there is a special case of Lemma 3.5. The finest tiling $\hat{1}$ operationalizes the Bramoullé-Kranton (2007) free-ride structure at the design level: each regulator free-rides on the tile immediately downstream. Elliott and Golub (2019) show that interior action profiles in network games are Pareto efficient only when the spectral radius of the benefits matrix equals one; since a forest is a directed acyclic graph, its benefits matrix is nilpotent and the spectral condition fails. The design problem escapes this trap by creating cooperative structure within tiles while keeping the inter-tile links acyclic.

On the fiscal side, the Galois connection gives a precise counterpart to Besley and Coate's (2003) finding that political economy frictions raise the effective centralization threshold above the Oates level. When politics makes map redesign unattractive, the grant dual provides an alternative that achieves the same welfare without redrawing borders; the political economy premium translates directly into fiscal cost. Galeotti, Golub, and Goyal (2020) study optimal targeting on a fixed graph; their eigenvector condition and our per-boundary threshold μ_v^* are complementary—GGG optimizes grants on the current map, while the Galois connection says when to change the map instead.

Outline. Section 2 builds the domain theory: forest domains, watershed tilings, and the Boolean algebra $\mathcal{W}$. Section 3 analyzes the provision game: equilibrium, welfare monotonicity, the grant-tiling duality, and the watershed envelope of a bad map. Section 4 extends the model to non-watershed partitions and to domains with local rather than transitive spillovers. Section 5 concludes with open empirical and theoretical questions.

2 Orderly Domains

Section 3 will need three things from the domain theory: a space of admissible maps, a way to move between maps, and a structural description of what those moves cost. This section builds all three. The payoff, stated upfront: the set of watershed tilings of a forest domain is a Boolean algebra $\mathcal{W}(V, \to) \cong B_{|V|-k}$ (Proposition 2.7). Drawing a jurisdictional map is equivalent to independently flipping $|V| - k$ binary switches. Each switch corresponds to a single non-root node; flipping it open splits one jurisdiction into two. There are no compatibility constraints: any combination of switches is a valid map. Section 3 will ask which combination maximizes welfare.

2.1 Domains

Definition 2.1. A *spillover domain* is a pair $(V, \to)$ where V is a nonempty set of *locations* and $\to$ is a preorder on V: a relation that is reflexive ($v \to v$ for all v) and transitive ($v \to w$ and $w \to x$ imply $v \to x$). We read $v \to w$ as "v spills into w."

Throughout this section $(V, \to)$ is a *forest domain*: a partial order (additionally antisymmetric) in which every element has at most one direct downstream neighbor. The canonical model is a river network. Locations are points along a system of waterways; $v \to w$ whenever a contaminant released at v reaches w by traveling downstream. No point on a river drains into two distinct channels—that is the forest condition.[1]

2.2 Watershed Tilings

A policymaker draws a map: a partition of V into tiles, one tile per jurisdiction. Not every domain admits a meaningful notion of "boundary placement." Working in a forest domain is not merely a simplifying assumption: by Proposition 2.5 below, admitting a watershed partition and being a forest domain are *equivalent conditions*. The forest assumption is thus the minimal hypothesis under which the partition theory we develop is possible at all. It is also sufficient: the set of watershed tilings of a forest domain is not merely a set but a Boolean algebra, with a clean combinatorial structure that will drive every result in Section 3.

Definition 2.2. A *tiling* of $(V, \to)$ is a partition $\mathcal{T} = \{T_1, \ldots, T_k\}$ of V. A tile $T \in \mathcal{T}$ is a *watershed tile* if the restriction of $\to$ to T has a unique maximal element—a point $d_T \in T$ such that $v \to d_T$ for all $v \in T$. We call d_T the *outlet* of T. A tiling is a *watershed tiling* if every tile is a watershed tile.

The outlet is determined by the tile and the domain—the mapmaker does not choose it. Drawing the boundaries determines where the flow exits each tile.

---

[1] The set V need not be finite or discrete. A river reach is naturally modeled as a one-dimensional manifold with an orientation that generates the preorder by downstream reachability. On a continuous domain G is infinite with no isolated points: no atom e_v exists and $\mathcal{W}(V, \to)$ is an atomless Boolean algebra, the continuous analogue of 2^G. The main text assumes V finite. Section 4.1 relaxes the forest assumption to domains with forks (deltas, distributaries); Section 4.2 drops transitivity, admitting cycles and mutual spillovers.

Example 2.3 (River system)**.** Let V consist of four rivers: the Allegheny, Monongahela, Ohio, and Mississippi, with spillover generated by

$$\text{Allegheny} \to \text{Ohio}, \quad \text{Monongahela} \to \text{Ohio}, \quad \text{Ohio} \to \text{Mississippi}.$$

The single-tile tiling $\{V\}$ is a watershed tile with outlet Mississippi. The tiling $\{T_1, T_2\}$ with $T_1 = \{\text{Allegheny, Monongahela, Ohio}\}$ and $T_2 = \{\text{Mississippi}\}$ gives two watershed tiles, with outlets Ohio and Mississippi respectively.

Example 2.4 (A non-watershed tile)**.** Let $V = \{A, B, C\}$ with $A \to B$ and $A \to C$ (a fork). The tile $T = \{A, B, C\}$ is not a watershed tile: A flows into both B and C, so there is no unique outlet. No tiling that groups A with both B and C in the same tile can be a watershed tiling.

When does a watershed tiling exist? A tile T has a unique outlet iff the restriction of $\to$ to T is a tree (a connected forest with one root). A tiling is a watershed tiling iff each tile is such a tree. This is possible iff $(V, \to)$ itself is a forest: the delta example shows that a fork prevents any tile containing the fork point from having a unique outlet.

Proposition 2.5. *$(V, \to)$ admits a watershed tiling if and only if it is a forest domain.*

Proof. ($\Rightarrow$) If some $v \in V$ has two distinct downstream neighbors w, w' with w and w' incomparable, then any tile containing v must contain a path toward both w and w', giving two candidates for the outlet. So no tile containing v can be a watershed tile.

($\Leftarrow$) If $(V, \to)$ is a forest, take the tiling whose tiles are the trees of the forest. Each tree has a unique root, which is the outlet. □

The poset of watershed tilings. Two watershed tilings are comparable by *refinement*: $\mathcal{T} \preceq \mathcal{T}'$ if $\mathcal{T}'$ is finer (each tile of $\mathcal{T}'$ is contained in a tile of $\mathcal{T}$). Write $\mathcal{W}(V, \to)$ for the poset of watershed tilings ordered by refinement.

The coarsest element of $\mathcal{W}(V, \to)$ takes each tree of the forest as a single tile. Finer tilings subdivide the trees by pushing outlets upstream. In a forest, the key structural fact is:

Lemma 2.6. *In a forest domain, the downset of any element is a chain: if $v \to w$ and $v \to x$, then w and x are comparable.*

Proof. In a forest, every element has a unique minimal element below it. If $v \to w$ and $v \to x$, both w and x lie on the unique downward path from v to the root, hence they are comparable. $\square$

This lemma is what makes watershed tilings well-behaved. It implies that within a forest, any two tiles that share a potential outlet must be nested—one refines the other. As a consequence:

Proposition 2.7. *For a forest domain with k trees and $|V|$ total nodes,*

$$\mathcal{W}(V, \to) \;\cong\; B_{|V|-k},$$

the Boolean algebra on $|V| - k$ generators.

Proof. Each tree contributes a Boolean algebra. Fix a tree T with n nodes and root r. A watershed tiling of T is determined by specifying, for each non-root node $v \in T \setminus \{r\}$, whether v serves as an outlet—that is, whether a tile boundary falls just upstream of v. These $n - 1$ choices are independent: whether v is an outlet constrains nothing about any other node u. The watershed tilings of T therefore correspond bijectively to subsets of $T \setminus \{r\}$, ordered by inclusion, and form the Boolean algebra B_{n-1}. The meet $\mathcal{T} \wedge \mathcal{T}'$ corresponds to intersection of outlet sets (the coarser tiling); the join to union (the finer tiling).

Trees are independent. Tilings of distinct trees share no nodes and place no constraints on one another, so

$$\mathcal{W}(V, \to) \;\cong\; \prod_i B_{n_i - 1} \;\cong\; B_{\sum_i (n_i - 1)} \;=\; B_{|V|-k},$$

using $B_m \times B_n \cong B_{m+n}$ and $\sum_i n_i = |V|$. $\square$

Theorem 2.8 (Birkhoff representation of $\mathcal{W}$)**.** *Let $R \subseteq V$ be the set of roots and $G = V \setminus R$ the non-root nodes. The map*

$$\phi : \mathcal{W}(V, \to) \;\longrightarrow\; 2^G, \qquad \phi(\mathcal{T}) \;=\; \{\, d_T : T \in \mathcal{T},\; d_T \in G \,\},$$

is an isomorphism of Boolean algebras, where 2^G is ordered by inclusion. Concretely:

- *(i)* (Surjectivity.) *Every $S \subseteq G$ is the outlet set of a unique watershed tiling $\phi^{-1}(S)$: assign each $v \in V$ to the tile whose outlet is the nearest element of $S \cup R$ downstream of v.*

(ii) (Order.) *Refinement corresponds to inclusion:* $\mathcal{T} \preceq \mathcal{T}'$ *iff* $\phi(\mathcal{T}) \subseteq \phi(\mathcal{T}')$.

(iii) (Atoms.) *The atoms of* $\mathcal{W}(V, \to)$ *are* $e_v := \phi^{-1}(\{v\})$ *for* $v \in G$. *They are mutually incomparable and form a discrete antichain of size* $|V| - k$.

Proof. Immediate from Proposition 2.7 and the bijection established in its proof. Surjectivity holds because any $S \subseteq G$ produces a valid watershed tiling: in a forest the downstream path from each v is unique, so the nearest element of $S \cup R$ downstream of v is well-defined, giving a unique outlet assignment. Atoms are incomparable because singleton outlet sets are pairwise non-nested. □

What the generators do. The atom e_v is the watershed tiling produced by placing exactly one non-root boundary, at v. In e_v, the tile of v consists of v and all nodes that flow into v before hitting any root (i.e., the upward cone of v within its tree); the rest of v's tree merges into a single tile with outlet the root r_v. Every other tree is left as a single tile.

Every watershed tiling is the join $\mathcal{T} = \bigvee_{v \in \phi(\mathcal{T})} e_v$: drawing a map means independently deciding, for each non-root node v, whether to activate the atomic boundary at v. There are no compatibility constraints—any combination is valid. The design problem in Section 3 is exactly: for which v should the switch be open?

Example 2.9 (Chain)**.** Let $(V, \to)$ be the chain $a \to b \to c$ (root c, so $G = \{a, b\}$). The atoms are $e_a = \{\{a\}, \{b, c\}\}$ and $e_b = \{\{a, b\}, \{c\}\}$. The full lattice B_2:

$$\begin{aligned}
\hat{0} &= \phi^{-1}(\emptyset) = \{a, b, c\}, \\
e_a &= \phi^{-1}(\{a\}) = \{a\},\ \{b, c\}, \\
e_b &= \phi^{-1}(\{b\}) = \{a, b\},\ \{c\}, \\
\hat{1} &= \phi^{-1}(\{a, b\}) = \{a\},\ \{b\},\ \{c\}.
\end{aligned}$$

The atoms e_a and e_b are incomparable: neither is a refinement of the other. Note that a is strictly upstream of b in $(V, \to)$, yet this order relation between the nodes is invisible in $\mathcal{W}$—both generators are free.

Example 2.10 (Fork)**.** Let $(V, \to)$ be the fork $a \to c$, $b \to c$ (root c, $G = \{a, b\}$). The atoms are $e_a = \{\{a\}, \{b, c\}\}$ and $e_b = \{\{a, c\}, \{b\}\}$, again giving B_2. Here a and b are already incomparable in $(V, \to)$, and they remain incomparable as generators—the same lattice arises from a structurally different domain.

The hypercube structure. The covering relations in $\mathcal{W}(V, \to)$ are simple: $\mathcal{T} \lessdot \mathcal{T}'$ if and only if $\phi(\mathcal{T}') = \phi(\mathcal{T}) \cup \{v\}$ for exactly one $v \in G$—one new boundary is activated, splitting exactly one tile in two. The Hasse diagram of $\mathcal{W}(V, \to)$ is therefore the $|V| - k$ dimensional hypercube $H_{|V|-k}$.

The meets and joins in $\mathcal{W}$—intersection and union of outlet sets—agree with those in the full partition lattice $\mathrm{Part}(V)$: the coarsest common coarsening and finest common refinement of any two watershed tilings are again watershed tilings. So $\mathcal{W}(V, \to)$ is a sublattice of $\mathrm{Part}(V)$.

What $\mathcal{W}$ sees and does not see. The isomorphism $\mathcal{W}(V, \to) \cong B_{|V|-k}$ depends on $(V, \to)$ only through $|V|$ and k. The chain $a \to b \to c$ and the fork $a \to c$, $b \to c$ give the same lattice B_2, though their tile geometries differ (Examples 2.9–2.10). The forest order on G is invisible to $\mathcal{W}$ as an abstract lattice, and the reason is precise: whether u is upstream or downstream of v is irrelevant to whether both can simultaneously serve as tile outlets. In a forest, they always can. The forest order governs *tile geometry*—which locations belong to the tile of outlet v—but not *outlet combinatorics*. The partition theory of a forest domain and its spillover theory are separable at this level: the lattice of admissible maps is the same for any two forest domains with the same $|V|$ and k, even if their internal structures differ entirely.

The preceding results describe the *space* of watershed tilings as an abstract combinatorial object. We now ask what a particular tiling *does* to the domain: how it transforms $(V, \to)$ into a coarser forest on fewer nodes.

2.3 The Quotient Domain

Once a watershed tiling $\mathcal{T}$ is chosen, each tile behaves as a single node in a coarser forest: the tile's outlet d_T is its point of contact with the rest of the domain, and inter-tile spillovers flow from outlet to outlet. The provision game in Section 3 is played on this coarser structure—the *quotient domain*—not on $(V, \to)$ directly. The key fact is that the quotient domain is itself a forest, so every result about forest domains applies recursively to the game.

Proposition 2.11 (Quotient domain)**.** *Let $\mathcal{T}$ be a watershed tiling of $(V, \to)$. Equip $\mathcal{T}$ with the* induced relation

$$T_i \to_{\mathcal{T}} T_j \iff d_{T_i} \to d_{T_j}.$$

Then:

(i) $(\mathcal{T}, \to_{\mathcal{T}})$ *is a forest domain.*

(ii) The tile-assignment map $\pi : (V, \to) \to (\mathcal{T}, \to_{\mathcal{T}})$, $\pi(v) = T(v)$, *is a morphism of forest domains—a functor between thin categories (Fong & Spivak, 2019; Davey & Priestley, 2002).*

We call $(\mathcal{T}, \to_{\mathcal{T}})$ *the* quotient domain *of* $(V, \to)$ *under* $\mathcal{T}$.

Proof. Reflexivity and transitivity of $\to_{\mathcal{T}}$ follow from those of $\to$. No forks: if $T_j \to_{\mathcal{T}} T_m$ and $T_j \to_{\mathcal{T}} T_n$, then $d_{T_j} \to d_{T_m}$ and $d_{T_j} \to d_{T_n}$; Lemma 2.6 gives d_{T_m} and d_{T_n} comparable, so the tile domain has no forks. Monotonicity of π: if $v \to w$ and $T(v) = T(w)$, then $\pi(v) = \pi(w)$ and monotonicity is reflexivity. If $v \to w$ and $T(v) \neq T(w)$, the path from v downstream exits $T(v)$ at $d_{T(v)}$ before reaching w, giving $d_{T(v)} \to d_{T(w)}$, hence $\pi(v) \to_{\mathcal{T}} \pi(w)$. $\square$

Corollary 2.12. $\mathcal{W}(\mathcal{T}, \to_{\mathcal{T}}) \cong {\downarrow}\mathcal{T}$ *in* $\mathcal{W}(V, \to)$*: the watershed tilings of the quotient domain are exactly the watershed tilings of* $(V, \to)$ *that are coarser than* $\mathcal{T}$.

Proof. A watershed tiling $\mathcal{T}'$ of $(\mathcal{T}, \to_{\mathcal{T}})$ groups tiles of $\mathcal{T}$ into supertiles, each with a unique outlet in the tile domain. The supertiles are unions of tiles of $\mathcal{T}$ whose outlets flow into a common tile-domain outlet, giving a coarser watershed tiling of $(V, \to)$ with $\phi(\cdot) \subseteq \phi(\mathcal{T})$. The correspondence $\mathcal{T}' \mapsto \phi(\mathcal{T}')$ is then a bijection onto $2^{\phi(\mathcal{T})} = {\downarrow}\mathcal{T}$ in $\mathcal{W}(V, \to) \cong 2^G$. $\square$

The watershed adjunction. The tile-assignment map π and the outlet map d are not merely inverse constructions—they form an adjoint pair, and the adjunction condition is precisely the watershed property. This is the structural core of Section 2: the watershed condition on tilings is not a constraint imposed from outside but the *definition* of a right adjoint to π.

The outlet map $d : (\mathcal{T}, \to_{\mathcal{T}}) \to (V, \to)$, $d(T) = d_T$, is a functor by definition of $\to_{\mathcal{T}}$, and it is a right adjoint to π:

$$\pi \dashv d, \qquad \pi(v) \leq_{\mathcal{T}} T \iff v \leq d(T).$$

($\Rightarrow$) If $\pi(v) \leq_{\mathcal{T}} T$ then $d_{T(v)} \leq d_T$, and $v \leq d_{T(v)} \leq d_T$. ($\Leftarrow$) If $v \leq d_T$, the path from v downstream exits $T(v)$ at $d_{T(v)}$ and then reaches d_T, giving $d_{T(v)} \leq d_T$, hence $\pi(v) \leq_{\mathcal{T}} T$.

The *unit*, $v \leq d(\pi(v)) = d_{T(v)}$, says every location flows downstream to its tile's outlet: this is the watershed property. The *counit*, $\pi(d(T)) = T(d_T) =$

$T \leq T$, is reflexivity. The watershed condition is therefore exactly the unit condition for $\pi \dashv d$: requiring a right adjoint to π is precisely the requirement that each tile drain to a unique outlet.

The poset $\mathcal{W}(V, \to)$ is the poset of *resolutions* of $(V, \to)$: each watershed tiling produces a quotient forest domain, and the refinement order on tilings corresponds to the refinement of resolutions. At the coarsest extreme (one tile per tree), the quotient domain is discrete—all inter-tile spillovers vanish. At the finest, the quotient approaches $(V, \to)$ itself.

Remark 2.13 (Single coarsening step)*.* A covering relation $\mathcal{T}' \lessdot \mathcal{T}$ in $\mathcal{W}(V, \to)$—one step toward coarser—corresponds to removing exactly one atom e_v from the outlet set: the tile with outlet v is absorbed into its unique downstream neighbor tile, reducing the tile count by one and collapsing the single inter-tile edge $T_v \to_{\mathcal{T}} T_{d(v)}$ in the quotient domain. Every coarsening decomposes into a sequence of such steps; every step is reversible. Because $\mathcal{W}(V, \to) \cong B_{|V|-k}$ is a Boolean algebra, these steps are independent: the welfare effect of merging at v does not depend on what has been merged elsewhere. This independence is what allows Section 3 to reduce the $2^{|V|-k}$ design problem to a single covering-step calculation, then read off the global optimum from the sign of that calculation.

3 Games on Watershed Tilings

Section 2 built the map: a forest domain $(V, \to)$ admits a Boolean lattice $\mathcal{W}(V, \to) \cong B_{|V|-k}$ of watershed tilings, each tiling producing a quotient forest on which jurisdictions interact through directional spillovers. This section asks what the map is *for*: given a forest domain, a spillover structure, and a provision game played on each tiling, which tiling should a planner choose?

The answer is both simple and surprising. *Watershed-aligned jurisdictions are optimal regardless of whether spillovers are beneficial or harmful.* The sign of the spillover parameter α determines which end of the watershed lattice is optimal—finest tiling $\hat{1}$ for $\alpha > 0$, coarsest tiling $\hat{0}$ for $\alpha < 0$ beyond a threshold—but in neither case does any non-watershed partition compete. The watershed lattice $\mathcal{W}(V, \to)$ is the right domain; the sign of α picks the pole.

Three further results deepen this picture. First, the two design instruments available to the planner—*map design* (which tiling to draw) and *grant design* (which fiscal transfers to issue on a given map)—are not separate problems but opposite faces of an anti-tone Galois connection mediated by the Boolean structure of $\mathcal{W}$. Second, any politically-drawn map $\mathcal{P}$ has a canonical watershed-aligned comparator $\mathcal{P}^{ws}$, constructed by projecting $\mathcal{P}$ onto $\mathcal{W}(V, \rightarrow)$ via a two-step closure operator. Third, the fiscal cost of maintaining $\mathcal{P}$ instead of $\mathcal{P}^{ws}$—what we call the *cost of bad borders*—is directly computable from the Galois connection. Section 3.6 connects these results to the literatures on fiscal federalism, network public goods, and optimal targeting.

The model is as follows. Fix a forest domain $(V, \rightarrow)$ and a watershed tiling $\mathcal{T} \in \mathcal{W}(V, \rightarrow)$. Each tile $T \in \mathcal{T}$ is assigned a single regulator who chooses a provision level $g_T \geq 0$. Provision is deployed at the tile's outlet d_T; it flows downstream to each tile directly adjacent in the quotient domain, arriving at a signed fraction $\alpha \in (-1, 1)$ of its source level. Positive α models a *benefit spillover*: upstream provision of a public good raises downstream quality (Bloch & Zenginobuz, 2007). Negative α models a *harm spillover*: upstream activity degrades downstream quality, the canonical case studied by Oates (1972). At $\alpha = 0$ tiles are independent.

3.1 The provision game

Definition 3.1 (Provision game on $\mathcal{T}$)**.** Fix $\mathcal{T} \in \mathcal{W}(V, \rightarrow)$ and $\alpha \in (-1, 1)$. The *provision game* $\Gamma(\mathcal{T}, \alpha)$ is the normal-form game in which:

(i) Each player $T \in \mathcal{T}$ chooses effort $g_T \geq 0$.

(ii) The *composite public good* received by T is

$$\gamma_T(\mathbf{g}) \; = \; g_T \; + \; \alpha \sum_{\substack{T' \in \mathcal{T} \\ T' \rightarrow_{\mathcal{T}} T}} g_{T'}.$$

(iii) Player T's payoff is

$$U_T(\mathbf{g}) \; = \; b(\gamma_T(\mathbf{g})) - g_T,$$

where $b : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is the benefit function.

The sum in (ii) runs over tiles that are *direct upstream neighbors* of T in the quotient domain $(\mathcal{T}, \rightarrow_{\mathcal{T}})$. Since $(\mathcal{T}, \rightarrow_{\mathcal{T}})$ is a forest (Proposition 1.4), each tile has at most finitely many such neighbors.

Assumption 1. b is twice continuously differentiable on $\mathbb{R}_{>0}$, strictly increasing, and strictly concave, with $b'(0) > 1$ and $\lim_{x\to\infty} b'(x) < 1$.

Assumption 1 ensures each player has a unique desired total consumption level $\bar{\gamma} = (b')^{-1}(1)$; provision is worthwhile from scratch but eventually too costly.

For $\alpha > 0$ the payoffs exhibit *strategic substitutes*: an increase in upstream provision shifts T's best reply downward. For $\alpha < 0$ (harm spillover) they exhibit *strategic complements*: upstream harm forces T to increase its own provision. The best-reply function for interior solutions is

$$\hat{g}_T(\mathbf{g}_{-T}) = (b')^{-1}(1) \;-\; \alpha \sum_{\substack{T'\in\mathcal{T}\\ T'\to_{\mathcal{T}} T}} g_{T'},$$

with the constraint $g_T \geq 0$. Let $\Delta(\mathcal{T}) = \max_{T\in\mathcal{T}} |\{T' : T' \to_{\mathcal{T}} T\}|$ denote the maximum in-degree of the quotient domain.

Proposition 3.2 (Existence and uniqueness)**.** *Under Assumption 1 and* $|\alpha|\,\Delta(\mathcal{T}) < 1$*, the provision game* $\Gamma(\mathcal{T},\alpha)$ *has a unique Nash equilibrium* $\mathbf{g}^*(\mathcal{T})$*.*

Proof. The strategy space $\prod_T [0,\bar{\gamma}]$ is compact and convex. Each best-reply map $\hat{g}_T$ is continuous and, under $|\alpha|\,\Delta(\mathcal{T}) < 1$, the joint best-reply map $\hat{\mathbf{g}} : \mathbf{g} \mapsto (\hat{g}_T(\mathbf{g}_{-T}))_T$ is a contraction in the ℓ^∞ norm:

$$\|\hat{\mathbf{g}}(\mathbf{g}) - \hat{\mathbf{g}}(\mathbf{g}')\|_\infty \;\leq\; |\alpha|\,\Delta(\mathcal{T})\,\|\mathbf{g}-\mathbf{g}'\|_\infty \;<\; \|\mathbf{g}-\mathbf{g}'\|_\infty.$$

Banach's fixed-point theorem gives a unique fixed point, which is the unique Nash equilibrium. □

Remark 3.3. Bloch and Zenginobuz (2007) establish uniqueness for $\alpha > 0$ under the same contraction condition; the formulation here extends their result to signed spillovers $\alpha \in (-1, 1)$ via $|\alpha|\,\Delta(\mathcal{T}) < 1$. For a chain quotient ($\Delta = 1$) the condition reduces to $|\alpha| < 1$, satisfied by hypothesis. The sign of α does not affect existence or uniqueness—but it determines everything about the welfare comparison. For $\alpha > 0$ the game exhibits strategic substitutes and finer tilings are better; for $\alpha < 0$ it exhibits strategic complements and coarser tilings may be better. Section 3.3 makes this precise.

3.2 Welfare and the design problem

The planner maximizes the utilitarian sum of equilibrium payoffs across all tiles.

Definition 3.4 (Equilibrium welfare)**.** For $\mathcal{T} \in \mathcal{W}(V, \to)$, the *equilibrium welfare* under tiling $\mathcal{T}$ is

$$W(\mathcal{T}) \;=\; \sum_{T \in \mathcal{T}} U_T(\mathbf{g}^*(\mathcal{T})) \;=\; \sum_{T \in \mathcal{T}} \Big[b(\gamma_T^*(\mathcal{T})) - g_T^*(\mathcal{T}) \Big].$$

The *tiling design problem* is to find $\mathcal{T}^* \in \arg\max_{\mathcal{T} \in \mathcal{W}(V,\to)} W(\mathcal{T})$. Because $\mathcal{W}(V, \to) \cong B_{|V|-k}$, this is a binary optimization over the non-root nodes $G = V \setminus R$: which $v \in G$ should be activated as tile outlets? A priori, this is a $2^{|G|}$-problem. The lattice structure will reduce it to a monotone problem on a totally ordered pair of poles—but only after we understand how W varies as we move through the lattice one step at a time.

3.3 Monotonicity in coarsening

The central claim of this section is that equilibrium welfare W is monotone on the watershed lattice $(\mathcal{W}(V, \to), \preceq)$, with direction determined by the sign of α. For $\alpha > 0$, finer is better: the unconstrained optimum is the finest tiling $\hat{1}$. For $\alpha < 0$ beyond a threshold α^*, coarser is better: the optimum is the coarsest tiling $\hat{0}$. In both cases the design problem collapses from a $2^{|G|}$-search to a choice between the two poles of the lattice.

Because $\mathcal{W}(V, \to) \cong B_{|V|-k}$ is a Boolean algebra, every coarsening is a sequence of single merges (Remark 2.13); it suffices to study what happens to W at a single covering step $\mathcal{T} \lessdot \mathcal{T}'$. The key structural fact about the NE is:

Lemma 3.5 (Composite efficiency)**.** *Under Assumption 1 and $|\alpha| \, \Delta(\mathcal{T}) < 1$, every tile $T \in \mathcal{T}$ achieves composite exactly $\bar{\gamma}$ at the Nash equilibrium: $\gamma_T^* = \bar{\gamma}$ for all T.*

Proof. The quotient domain $(\mathcal{T}, \to_{\mathcal{T}})$ is a forest, so its tiles admit a topological order from upstream to downstream. Process tiles in that order. A tile with no upstream neighbors sets $g_T^* = \bar{\gamma}$ by the interior first-order condition $b'(\gamma_T) = 1$, so $\gamma_T^* = \bar{\gamma}$. For any other tile T, the upstream neighbors have already been processed and satisfy $g_{T'}^* \geq 0$. For $\alpha > 0$, the contraction condition $|\alpha| \, \Delta < 1$ ensures $\alpha \sum_{T' \to T} g_{T'}^* < \bar{\gamma}$, so $g_T^* = \bar{\gamma} - \alpha \sum_{T' \to T} g_{T'}^* > 0$ and $\gamma_T^* = \bar{\gamma}$. For

$\alpha < 0$, every term in the sum enters with a positive sign ($-\alpha > 0$), so $g_T^* > 0$ and $\gamma_T^* = \bar{\gamma}$. □

Lemma 3.5 collapses the welfare comparison to a comparison of total provision effort.

Corollary 3.6. *$W(\mathcal{T}) = |\mathcal{T}|\, b(\bar{\gamma}) - G(\mathcal{T})$ where $G(\mathcal{T}) := \sum_{T\in\mathcal{T}} g_T^*$. Moreover,*

$$G(\mathcal{T}) \;=\; \frac{|\mathcal{T}|\,\bar{\gamma} + \alpha \sum_{r\in R_{\mathcal{T}}} g_r^*}{1+\alpha}, \tag{1}$$

where $R_{\mathcal{T}} \subseteq \mathcal{T}$ is the set of root tiles (those with no downstream neighbor in the quotient domain).

Proof. The welfare formula is immediate from Lemma 3.5. For (1): sum the equilibrium conditions $g_T^* = \bar{\gamma} - \alpha \sum_{T'\to T} g_{T'}^*$ over all T. Each non-root tile T' contributes $g_{T'}^*$ exactly once to the right-hand sum (to its unique downstream neighbor), so

$$G = |\mathcal{T}|\,\bar{\gamma} - \alpha\big(G - \textstyle\sum_r g_r^*\big).$$

Solving gives (1). □

For a covering step $\mathcal{T} \lessdot \mathcal{T}'$ (adding one tile), the welfare change is

$$W(\mathcal{T}') - W(\mathcal{T}) \;=\; b(\bar{\gamma}) - \Delta G, \qquad \Delta G := G(\mathcal{T}') - G(\mathcal{T}).$$

From (1), $\Delta G = (\bar{\gamma} + \alpha\, \Delta g_{r_0})/(1+\alpha)$, where r_0 is the unique root tile downstream of the split and $\Delta g_{r_0} = g_{r_0}^{*\prime} - g_{r_0}^*$.

Proposition 3.7 (Monotonicity, $\alpha > 0$)**.** *If $\alpha > 0$, then for every covering step $\mathcal{T} \lessdot \mathcal{T}'$,*

$$W(\mathcal{T}') \;>\; W(\mathcal{T}).$$

Equilibrium welfare is strictly increasing in refinement; the unconstrained optimal tiling is $\hat{1}$.

Proof. For $\alpha > 0$, Lemma 3.5 gives $g_{r_0}^{*\prime} = \bar{\gamma} - \alpha \sum_{T'\to r_0} g_{T'}^{*\prime} \le \bar{\gamma}$ (since all provisions are non-negative). Also $g_{r_0}^* \ge 0$. Hence $\Delta g_{r_0} \le \bar{\gamma}$, and

$$\Delta G \;\le\; \frac{\bar{\gamma} + \alpha\bar{\gamma}}{1+\alpha} \;=\; \bar{\gamma} \;<\; b(\bar{\gamma}).$$

Thus $W(\mathcal{T}') - W(\mathcal{T}) = b(\bar{\gamma}) - \Delta G \ge b(\bar{\gamma}) - \bar{\gamma} > 0$. □

Economic interpretation. Splitting tile A at boundary v creates a downstream tile D that achieves composite $\bar{\gamma}$ at cost $g_D^* = \bar{\gamma} - \alpha\, g_U^* < \bar{\gamma}$—it free-rides on upstream provision. The new tile delivers $b(\bar{\gamma})$ in benefits at below-$b(\bar{\gamma})$ cost. Since $b(\bar{\gamma}) > g_D^*$ (strictly), welfare strictly rises. The finest tiling $\hat{1}$ maximises the number of such free-ride edges.

The $\alpha < 0$ case. For $\alpha < 0$, Lemma 3.5 still holds but now $g_{r_0}^{*'} \geq \bar{\gamma}$ (the root must over-provide to offset upstream harm), so $\Delta g_{r_0} \geq -g_{r_0}^*$ can be positive and large. Coarser is better iff $\Delta G > b(\bar{\gamma})$, i.e.,

$$\bar{\gamma} + \alpha\, \Delta g_{r_0} \;>\; (1+\alpha)\, b(\bar{\gamma}).$$

In the two-tile chain, $\Delta g_{r_0} = |\alpha|\bar{\gamma}$, giving $W(\hat{1}) < W(\hat{0})$ iff

$$|\alpha| \;>\; \alpha^* \;:=\; \frac{b(\bar{\gamma}) - \bar{\gamma}}{\bar{\gamma}}. \tag{2}$$

The threshold α^* is the bliss-point surplus per unit provision. Internalization dominates when the spillover magnitude exceeds this surplus. Whether $\alpha^* < 1$—i.e., whether the Oates regime exists within the model's parameter space—depends on the benefit function: $\alpha^* < 1$ iff $b(\bar{\gamma}) < 2\bar{\gamma}$.

Example 3.8 (Sign is not enough)**.** Let $b(x) = 2x - x^2$, so $\bar{\gamma} = \frac{1}{2}$, $b(\bar{\gamma}) = \frac{3}{4}$, and $\alpha^* = \frac{1}{2}$. Consider the two-node chain $a \to b$ and compare $\hat{0} = \{\{a, b\}\}$ with $\hat{1} = \{\{a\}, \{b\}\}$:

α	g_a^*	g_b^*	$W(\hat{1})$	$W(\hat{0})$
-0.3	$\frac{1}{2}$	0.65	0.35	0.25
-0.5	$\frac{1}{2}$	0.75	0.25	0.25
-0.7	$\frac{1}{2}$	0.85	0.15	0.25

At $\alpha = -0.3$ the spillover is negative—upstream harms downstream, forcing b to over-provide by 0.15—yet finer is still better, because the over-provision cost ($|\alpha|\bar{\gamma} = 0.15$) is less than the benefit premium ($b(\bar{\gamma}) - \bar{\gamma} = 0.25$). At $\alpha = -0.7$ the harm is large enough to dominate and coarser wins. The crossover is exactly at $\alpha = -\alpha^* = -\frac{1}{2}$.

Remark 3.9. Both mechanisms privilege watershed-aligned borders, but for opposite reasons. For $\alpha > 0$, non-watershed borders destroy free-ride edges (grouping a provider with its beneficiary collapses the rent). For $\alpha < 0$ with

$|\alpha| > \alpha^*$, non-watershed borders prevent harm internalization (separating a polluter from the territory it harms). In either case restricting the design problem to $\mathcal{W}(V, \to)$ is without loss; the sign of α determines which pole of the lattice is optimal.

Remark 3.10. The constrained design problem—find the optimal tiling given a fixed number m of tiles—is well-posed in both regimes via the Boolean structure of $\mathcal{W}(V, \to) \cong B_{|V|-k}$. Under $\alpha > 0$, the unconstrained optimum is $\hat{1}$ and constraints force coarsening; the question is which of the $|V| - k - m$ boundaries to remove at least welfare cost. Under $\alpha < 0$ with $|\alpha| > \alpha^*$, the unconstrained optimum is $\hat{0}$ and constraints prevent full merger; the question is which $m - k$ boundaries to retain. In both cases the ranking criterion is $b(\bar{\gamma}) - \Delta G_v$: the net welfare gain from activating boundary v. Section 3.4 shows this is also the quantity that determines the break-even fiscal grant for v, connecting the constrained map-design problem to the grant-design problem via a single common statistic.

3.4 The grant-tiling duality

The map design problem of §3.3 assumed the planner could freely choose any watershed tiling. In practice, jurisdictional maps are drawn by legislatures, reflect historical boundaries, and change slowly if at all. A planner who cannot redraw the map faces a second question: given a fixed partition $\mathcal{P}$, which intergovernmental grants can replicate the welfare outcome of a better-aligned map?

The mechanism is straightforward. When a watershed boundary falls across two jurisdictions, the upstream regulator does not fully internalize the enforcement consequences that fall downstream; the downstream regulator compensates by over-providing. A grant to the upstream tile at each misaligned boundary corrects this: it pays the upstream regulator to internalize what it is externalizing, shifting its effective incentive to act as though the boundary did not exist. Calibrated to the break-even level μ_v^*, the grant replicates the merged outcome without merging.

The main result is that this equivalence is exact and runs both ways. *Map design* and *grant design* are not merely alternatives; they are dual instruments connected by an anti-tone Galois connection on the Boolean algebra $\mathcal{W}(V, \to)$. For any tiling, the minimum grant schedule that supports it is computable in closed form; for any grant schedule, the coarsest tiling it makes welfare-optimal is equally computable. The duality makes the *cost of bad borders* precise: not a welfare bound but an exact fiscal sum, one term μ_v^* per misaligned boundary, indexed by the same Boolean algebra that parametrizes the map.

Grant-augmented game. A *grant schedule* $\mu \in \mathbb{R}^G$ assigns to each non-root node $v \in G$ a transfer μ_v paid (by the planner) to tile $T_{d(v)}$ per unit provided by tile T_v, augmenting the natural spillover on active boundary v to an effective value $\tilde{\alpha}_v := \alpha + \mu_v$. The grant schedule lives on $G = V \setminus R$, indexed by the same set as the Boolean algebra $\mathcal{W}(V, \to) \cong 2^G$—and is therefore tiling-independent: the same space $M := \mathbb{R}^G$ serves regardless of which tiling is in force.

Under tiling $\mathcal{T}$ and grant μ, grant μ_v is active on edge $T_v \to_{\mathcal{T}} T_{d(v)}$ only when $v \in \phi(\mathcal{T})$ (i.e., when the boundary is open). When $v \notin \phi(\mathcal{T})$ the edge does not exist in the quotient domain and μ_v is irrelevant. Lemma 3.5 extends immediately to $\Gamma(\mathcal{T}, \alpha, \mu)$: composite efficiency holds for each active edge provided $|\tilde{\alpha}_v|\,\Delta(\mathcal{T}) < 1$, and the welfare formula $W(\mathcal{T}, \mu) = |\mathcal{T}|\,b(\bar{\gamma}) - G(\mathcal{T}, \mu)$ continues to hold with $\tilde{\alpha}_v$ in place of α on each edge.

Per-boundary break-even grant. For each non-root node v, the welfare gain from activating boundary v (splitting the tile at v) with effective spillover $\tilde{\alpha}_v$ is

$$\Delta W_v(\tilde{\alpha}_v) \;=\; b(\bar{\gamma}) - \Delta G_v(\tilde{\alpha}_v),$$

where ΔG_v is the change in total provision from the split (Corollary 3.6). Since ΔW_v is continuous and strictly increasing in $\tilde{\alpha}_v$ (Proposition 3.7 applied per edge), there is a unique threshold $\tilde{\alpha}_v^{thr}$ at which $\Delta W_v = 0$. The *break-even grant* for boundary v is

$$\mu_v^* \;:=\; \tilde{\alpha}_v^{thr} - \alpha. \tag{3}$$

Activating boundary v raises welfare iff $\tilde{\alpha}_v > \tilde{\alpha}_v^{thr}$, i.e., iff the grant exceeds μ_v^*.

For the two-tile chain, $\Delta G_v(\tilde{\alpha}) = \bar{\gamma}(1 - \tilde{\alpha})$, giving $\tilde{\alpha}_v^{thr} = -\alpha^*$ and

$$\mu_v^* \;=\; |\alpha| - \alpha^*.$$

This is positive in the Oates regime ($|\alpha| > \alpha^*$) and non-positive otherwise. The threshold α^* from (2) is precisely the break-even grant for the smallest forest.

The Galois connection. Order grant schedules component-wise: $\mu \leq \mu'$ iff $\mu_v \leq \mu'_v$ for all $v \in G$. Define the *grant-supported tiling*

$$\mathcal{T}^*(\mu) \;:=\; \phi^{-1}\big(\{v \in G : \mu_v \geq \mu_v^*\}\big),$$

the tiling that activates exactly the boundaries whose break-even grant is met. The break-even grant map is $\mu^* : \mathcal{W}(V, \to) \to M$, $\mu^*(\mathcal{T})_v = \mu_v^* \cdot \mathbf{1}[v \in \phi(\mathcal{T})]$: the schedule that funds precisely the active boundaries of $\mathcal{T}$ at their break-even levels.

Proposition 3.11 (Galois connection)**.** *The maps*

$$\mu^* : (\mathcal{W}(V, \to), \preceq) \to (M, \leq) \quad \textit{and} \quad \mathcal{T}^* : (M, \leq) \to (\mathcal{W}(V, \to), \preceq)$$

form an anti-tone Galois connection:

$$\mu^*(\mathcal{T}) \;\leq\; \mu \quad \Longleftrightarrow \quad \mathcal{T} \;\preceq\; \mathcal{T}^*(\mu).$$

Proof. Both directions follow from the definitions and the independence of boundary decisions (since $\mathcal{W} \cong 2^G$, there are no cross-boundary constraints).

($\Rightarrow$): If $\mu^*(\mathcal{T}) \leq \mu$ component-wise, then for every $v \in \phi(\mathcal{T})$, $\mu_v \geq \mu_v^*$, so $v \in \phi(\mathcal{T}^*(\mu))$. Hence $\phi(\mathcal{T}) \subseteq \phi(\mathcal{T}^*(\mu))$, i.e., $\mathcal{T} \preceq \mathcal{T}^*(\mu)$.

($\Leftarrow$): If $\mathcal{T} \preceq \mathcal{T}^*(\mu)$, then $\phi(\mathcal{T}) \subseteq \phi(\mathcal{T}^*(\mu)) = \{v : \mu_v \geq \mu_v^*\}$, so $\mu_v \geq \mu_v^*$ for all $v \in \phi(\mathcal{T})$, giving $\mu^*(\mathcal{T}) \leq \mu$. □

Remark 3.12 (What the Boolean structure buys). The proof uses $\mathcal{W} \cong 2^G$ in a precise way: the absence of cross-boundary constraints in the Boolean algebra means each boundary's activation decision is independent, so the Galois connection decomposes into $|G|$ independent one-dimensional problems. In a non-Boolean tiling lattice (e.g., the partition lattice of a non-forest domain), boundaries interact and the decomposition fails: there is no clean per-boundary break-even grant and no Galois connection of this form. The watershed condition—equivalently, the forest condition—is what enables the duality.

Remark 3.13 (The hidden problem). The standard analysis of jurisdiction design asks which map to draw (Problem A: choose $\mathcal{T}^* \in \mathcal{W}$). The dual problem asks which grants to design for a given map (Problem B: choose $\mu^*(\mathcal{T}) \in M$). Without the lattice structure, these look like two separate optimization problems. The Galois connection says they are opposite faces of the same object: $\mathcal{T}$ is sufficient without grants iff the zero grant supports $\mathcal{T}$, i.e., iff $\mu_v^* \leq 0$ for all $v \in \phi(\mathcal{T})$. The threshold α^* marks the boundary: for $|\alpha| \leq \alpha^*$, all boundaries are self-justifying ($\mu_v^* \leq 0$) and the optimal tiling is $\hat{1}$ with no grants needed; for $|\alpha| > \alpha^*$, some boundaries require positive grants to be welfare-improving, and the optimal grant-free tiling shifts toward $\hat{0}$.

3.5 The watershed envelope of a bad map

Political maps are drawn by history, not hydrology. States follow mountain ranges, rivers, and colonial borders that bear no relation to the catchment structure of $(V, \to)$. Given such a map—an arbitrary partition $\mathcal{P}$ of V—the theory developed in §§3.3–3.4 does not directly apply, because $\mathcal{P}$ may not be a watershed tiling. This section constructs the *watershed envelope* $\mathcal{P}^{ws} \in \mathcal{W}(V, \to)$—the canonical watershed-aligned comparator—and uses the Galois connection to give a precise account of the cost of bad borders.

Construction. Let $\{T_j\}_{j \in J}$ be the trees of the forest (connected components of $(V, \to)$ with roots $\{r_j\}$). Given $\mathcal{P}$, proceed in two steps.

Step 1 (Tree-alignment). For each tile $P_i \in \mathcal{P}$ and each tree T_j, form the piece $P_i^j := P_i \cap V(T_j)$. The resulting collection $\mathcal{P}^{(1)} := \{P_i^j : P_i^j \neq \emptyset\}$ is a partition of V that respects tree boundaries. Any tile of $\mathcal{P}$ that crosses a tree boundary is split; no cross-tree merging remains.

Step 2 (Watershed-coarsening). Each piece $Q \in \mathcal{P}^{(1)}$ lies within one tree T_j. Let $d(Q)$ be the *lowest common outlet* of Q: the unique most-downstream node $v \in V(T_j)$ such that every element of Q drains through v. (In a tree, the set $\{v : u \to^* v \text{ for all } u \in Q\}$ is a chain with a unique minimum element, namely $d(Q)$.) Define the *watershed envelope tile* $\mathcal{T}(d(Q)) := \{u \in V(T_j) : u \to^* d(Q)\}$, the upstream cone of $d(Q)$ within T_j. Set

$$\mathcal{P}^{ws} \;:=\; \text{the partition generated by } \big\{\, \mathcal{T}(d(Q)) : Q \in \mathcal{P}^{(1)} \big\},$$

obtained by merging any tiles whose watershed envelopes coincide.

Lemma 3.14. *$\mathcal{P}^{ws} \in \mathcal{W}(V, \to)$ for every partition $\mathcal{P}$ of V.*

Proof. Each tile of $\mathcal{P}^{ws}$ is an upstream cone $\mathcal{T}(d(Q))$ within one tree. An upstream cone has a unique outlet ($d(Q)$) and is closed under upstream neighbors: if $u \in \mathcal{T}(d(Q))$ and $w \to u$ then $w \to^* d(Q)$, so $w \in \mathcal{T}(d(Q))$. These are precisely the watershed tile conditions. Since the cones within each tree partition $V(T_j)$ (every node drains through a unique outlet in $\mathcal{P}^{ws}$), the collection is a valid tiling. $\square$

Two features of the construction are worth noting. First, $\mathcal{P}^{ws}$ is always *coarser* than $\mathcal{P}^{(1)}$: each piece Q is a subset of its watershed envelope tile $\mathcal{T}(d(Q))$. Second, the two-step projection is monotone in $\mathcal{P}$: a finer political map yields a finer watershed envelope.

The welfare comparison. Let $W(\mathcal{P})$ denote equilibrium welfare when the game is played on partition $\mathcal{P}$ (whether or not $\mathcal{P} \in \mathcal{W}$). For $\mathcal{P} \in \mathcal{W}$, $W(\mathcal{P})$ is the quantity analyzed in §3.3. Extending W to non-watershed partitions requires modeling the provision game on tiles without unique outlets—a tile with two exits has ambiguous spillover routing—and this extension is developed in §4.1.

Taking that extension as given, Proposition 3.15 is proved in §4.1 using the covering-step formula of §3.3 applied to the extended game.

The intuition runs through the two construction steps. Nodes in the non-outlet component of a cross-tree tile—those in a tree other than the outlet's tree—generate no inter-tile edges under the single-outlet convention, so they are invisible to the provision game. Their tree-alignment and subsequent watershed-coarsening leave every tree's game structure unchanged, contributing zero net welfare change. Within-tree non-watershed edges E_{orig} do affect the game; each is eliminated by a covering merge in Step 2. In the within-tree inter-tile chain of $n \geq 2$ tiles containing that edge, every tile at or below the merge point sheds provision: $\Delta G = \bar{\gamma}(1 + |\alpha| + \cdots + |\alpha|^{n-1}) \geq \bar{\gamma}(1 + |\alpha|) > b(\bar{\gamma}) = \bar{\gamma}(1 + \alpha^*)$ when $|\alpha| > \alpha^*$, so each merge raises welfare. The subtlety is that the two steps are not individually monotone for $\alpha < 0$: tree-alignment creates new within-tree inter-tile edges that force over-provision before Step 2 removes them. The proof therefore tracks only E_{orig}—the edges present in $\mathcal{P}$ before tree-alignment—and shows each contributes a net welfare gain; see Remark 4.2.

Proposition 3.15 (Watershed dominance)**.** *Under Assumption 1 and* $|\alpha|\,\Delta < 1$*, if* $|\alpha| > \alpha^*$ *then*

$$W(\mathcal{P}^{ws}) \;\geq\; W(\mathcal{P})$$

for every partition $\mathcal{P}$ *of* V*. Restricting the design problem to* $\mathcal{W}(V, \rightarrow)$ *is without loss whenever the spillover magnitude exceeds* α^**.*

The case $|\alpha| \leq \alpha^*$ (which includes all of $\alpha > 0$) runs in the opposite direction on Step 2: within-tree refinement raises welfare, so the relevant dominator is the finest watershed tiling $\hat{1}$, not $\mathcal{P}^{ws}$. In that regime the policy recommendation is: refine to $\hat{1}$ and issue no grants.

The cost of bad borders. Even without Proposition 3.15, the Galois connection of §3.4 gives a precise fiscal measure of the gap between any map $\mathcal{P}$ and its watershed comparator.

Corollary 3.16. *For any partition $\mathcal{P}$, the grant schedule $\mu^*(\mathcal{P}^{ws})$—the break-even grants for the boundaries that $\mathcal{P}^{ws}$ activates but $\mathcal{P}$ does not enforce—is the minimum fiscal transfer that makes $\mathcal{P}$ reproduce the welfare outcome of $\mathcal{P}^{ws}$. Its total cost $\|\mu^*(\mathcal{P}^{ws})\|_1 = \sum_{v \in \phi(\mathcal{P}^{ws}) \setminus \phi(\mathcal{P})} \mu_v^*$ is the* cost of bad borders*: the price, in fiscal transfers, of the mismatch between the political map and the watershed structure.*

Proof. By Proposition 3.11, the minimum grant that supports $\mathcal{P}^{ws}$ is $\mu^*(\mathcal{P}^{ws})$. A grant μ achieves the welfare outcome of $\mathcal{P}^{ws}$ on map $\mathcal{P}$ iff the effective spillover $\tilde{\alpha}_v = \alpha + \mu_v$ at each boundary $v \in \phi(\mathcal{P}^{ws}) \setminus \phi(\mathcal{P})$ satisfies $\tilde{\alpha}_v \geq \tilde{\alpha}_v^{thr}$, i.e., $\mu_v \geq \mu_v^*$. The minimum such schedule is $\mu_v = \mu_v^*$. ☐

Remark 3.17 (Oates's observation, formalized). Oates observed that "rivers are the worst sort of border": drawing a jurisdictional line along a waterway puts the externality-generating activity (upstream provision) and the externality-receiving territory (downstream) in different jurisdictions, preventing internalization. Proposition 3.15 quantifies this: for any partition $\mathcal{P}$, the welfare gap $W(\mathcal{P}^{ws}) - W(\mathcal{P}) \geq 0$ is the welfare cost of misalignment, and Corollary 3.16 gives its fiscal dual—the grant cost of maintaining the bad border. Together they provide a complete regulatory design recommendation from any historically-given map: either redraw to $\mathcal{P}^{ws}$, or fund the grants $\mu^*(\mathcal{P}^{ws})$.

3.6 Discussion

Section 3 set out to answer a design question: given a river network and a regulatory game played on its jurisdictions, which map should a planner draw? The answer that emerged is cleaner than we had a right to expect. This discussion unpacks the logic, connects it to the literature, and tries to make visible what the mathematical structure actually did.

One answer, two mechanisms. The headline result is that watershed-aligned jurisdictions are optimal regardless of the sign of α. For $\alpha > 0$—upstream provision benefits downstream, as in regional public goods or drinking-water quality—the finest watershed tiling $\hat{1}$ (one jurisdiction per node) is optimal. For $\alpha < 0$ with $|\alpha| > \alpha^*$—upstream activity harms downstream, as in pollution or nutrient runoff—the coarsest watershed tiling $\hat{0}$ (one jurisdiction per whole catchment) is optimal. At the boundary $|\alpha| = \alpha^*$ the two regimes meet: every tiling achieves the same welfare.

The two mechanisms are qualitatively different. Under positive spillovers, the gain from splitting a jurisdiction is a *free-ride dividend*: the downstream jurisdiction achieves composite $\bar{\gamma}$ at strictly less than full cost, because it inherits α times the upstream jurisdiction's provision for free. Each open boundary in $\phi(\mathcal{T})$ is a gift channel from upstream to downstream. Closing a boundary—merging two jurisdictions—destroys one such channel: the merged jurisdiction must now self-finance what it used to receive. More boundaries, hence finer tilings, are always better.

Under negative spillovers, the mechanism inverts. Upstream activity now imposes a cost on downstream: the downstream jurisdiction must over-provide by $|\alpha|$ times the upstream provision to reach $\bar{\gamma}$. This over-provision is pure waste from the merged planner's perspective. When $|\alpha| > \alpha^*$, the waste exceeds the bliss-point surplus $b(\bar{\gamma}) - \bar{\gamma}$, and closing the boundary—merger, internalization—is worth it. Fewer boundaries, coarser tilings, are better. Example 3.8 showed that sign alone does not determine direction: the threshold $\alpha^* = (b(\bar{\gamma}) - \bar{\gamma})/\bar{\gamma}$ is the *bliss-point surplus ratio*, a property of the benefit function that measures how much a jurisdiction values the composite above its marginal cost. Only when the harm magnitude clears this bar does internalization pay.

The Oates mechanism, formalized. Oates (1972) observed that jurisdictions separated by rivers suffer precisely because the river—and the regulatory boundary following it—cuts across the spillover flow, leaving each jurisdiction unable to account for its effects on the other. The canonical case is nutrient pollution: agricultural runoff from Pennsylvania farms flows down the Susquehanna, crosses into Maryland, and degrades the Chesapeake Bay. Maryland regulators cannot compel Pennsylvania farmers; Pennsylvania regulators have no incentive to internalize harms that fall on Maryland. The Chesapeake Bay Program, established in 1983, is the federal response: a voluntary compact among six states and the District of Columbia that functions, in effect, as a partial merger of the relevant jurisdictions along watershed lines. It is, in the language of this paper, a step from $\hat{1}$ toward $\hat{0}$ along the watershed lattice.

The present framework gives Oates's intuition three things it previously lacked: *precision*, *conditions*, and *a dual.*

Precision. The welfare gap $W(\mathcal{P}^{ws}) - W(\mathcal{P})$ between the watershed comparator and any political map $\mathcal{P}$ is not merely positive; it has a closed form via Corollary 3.6. For the two-node chain it equals $(|\alpha| - \alpha^*)\bar{\gamma}/(1+\alpha)$: zero at the threshold, growing linearly in the excess harm $|\alpha| - \alpha^*$. The gap is computable from observable parameters—the benefit function, the spillover magnitude, the tree structure—without any unobservable.

Conditions. The Oates result is not universal: it requires $|\alpha| > \alpha^*$. This condition fails when the spillover is weak or when jurisdictions have high bliss-point surplus (flat benefit functions far from the optimum). In such cases, even negative spillovers are better handled by the finest tiling. The statement "rivers are the worst sort of border" is correct in the Oates regime and *wrong* outside it. The lattice tells you which regime you are in.

A dual. The Galois connection of §3.4 says that for any map $\mathcal{P}$, there is an equivalent fiscal instrument: the grant schedule $\mu^*(\mathcal{P}^{ws})$ that makes $\mathcal{P}$ perform as well as the watershed comparator. This is not just a theoretical device. It is the Chesapeake Bay Program's operational reality: the program works primarily through intergovernmental transfers and matching grants (nutrient trading credits, agricultural best-practice subsidies) rather than by redrawing state lines. The duality says these two instruments—redraw the map vs. design the grants—are formally equivalent, faces of the same Galois connection. The cost of bad borders is then the program's minimum fiscal budget.

The political economy literature complicates the map-design side of this equivalence in important ways. Besley and Coate (2003) show that in a model where districts elect representatives and a legislature allocates shared spending, the case for centralization is substantially weaker than the standard spillover analysis implies. Political economy frictions—strategic delegation by voters, coalition formation in the legislature, uncertainty over budget shares—impose additional costs on the centralized regime that the benevolent-planner analysis ignores. Their key result is that the threshold spillover level above which centralization dominates is *strictly higher* than the standard Oates threshold, and may be outside the feasible range altogether when districts are sufficiently heterogeneous.

The present framework provides a precise economic benchmark (α^*) while deferring the political economy frictions. The Galois connection offers a different way to interpret the Besley-Coate tension: whenever political economy makes map redesign unattractive (because the threshold is higher than α^* once frictions are accounted for), the fiscal dual provides an alternative. The grant schedule $\mu^*(\mathcal{P}^{ws})$ achieves the same welfare as the watershed-aligned map without requiring the political act of redrawing borders. The program cost of this strategy is the Besley-Coate "political economy premium" translated into fiscal terms.

The BZ mechanism and its place in the literature. Bloch and Zenginobuz (2007) study a game of local public good provision on a network where each player benefits from neighbors' provision. In their setting every player free-rides on every neighbor simultaneously; the equilibrium features "convergence to the bottom" as players undercut each other. The present setup inherits this structure on each watershed tiling, but adds a design dimension: the planner chooses which neighborhood structure (quotient domain) to impose by drawing the map.

The free-ride dividend identified above is the constructive face of the BZ phenomenon. BZ treat free-riding as a coordination failure to be minimized; here it is a mechanism to be harnessed. Under $\alpha > 0$, splitting a jurisdiction extracts a free-ride gain that the merged planner would have internalized away. Put differently: merger is the wrong response to a positive spillover. The correct response is fragmentation—draw many small jurisdictions aligned with the flow, and let the downstream ones free-ride on the upstream ones. The finest watershed tiling $\hat{1}$ achieves this maximally.

This reversal of the usual instinct ("merger solves the externality problem") is possible only because provision is a strategic substitute under $\alpha > 0$: each regulator reduces effort when upstream provision increases. The equilibrium composite $\bar{\gamma}$ is achieved at collectively lower cost when the strategic substitution runs through open boundaries. The coarser the tiling, the more substitution is internalized and the more costly it is to reach $\bar{\gamma}$.

Tiebout (1956) argued that local competition among many small jurisdictions disciplines providers and achieves efficient sorting. The present result provides a complementary foundation for small jurisdictions that does not rest on mobility or sorting: even with fixed populations and no competition, small watershed-aligned jurisdictions can dominate a single large one when spillovers are beneficial. The mechanism is not competitive pressure but the geometry of the network.

The spectrum of the public goods literature. Situating the present game in the broader network public goods literature clarifies what is new. Bergstrom, Blume, and Varian (1986) study private provision of a *pure* public good: every player benefits from every other player's provision, and the unique Nash equilibrium features neutrality—small income transfers among contributors leave aggregate provision unchanged. This is exactly the game played on the coarsest tiling $\hat{0}$: a single tile is a BBV game. The Constant Composite Lemma (Lemma 3.5) is the directed-forest analogue of BBV neutrality: the aggregate composite $\bar{\gamma}$ is invariant across tilings, even as who bears the cost changes.

At the other extreme, Bramoullé and Kranton (2007) study the corner-solution version of the network game (linear benefit, constant marginal cost) and show that Nash equilibria correspond to *maximal independent sets* of the game graph. In a directed tree with edges pointing downstream (upstream to downstream), an independent set is a set of nodes with no two connected by an edge: no node in the set is the direct downstream neighbor of another. In a chain, the two maximal independent sets are the even-depth and odd-depth layers. The BK structure is the extreme version of the free-ride pattern identified above: exactly half the nodes exert full effort and the other half free-ride completely. The $\hat{1}$ tiling of our model operationalizes this at the design level—by drawing boundaries at every node, the planner ensures each tile's downstream neighbor is a distinct jurisdiction that free-rides on it.

Why design matters: the acyclicity observation. Elliott and Golub (2019) provide a striking characterization of Pareto efficiency in network games: an interior action profile is Pareto efficient if and only if the spectral radius of the benefits matrix $\mathbf{B}(\mathbf{a})$ equals one. A forest is a directed acyclic graph; its benefits matrix is nilpotent, so $r(\mathbf{B}(\mathbf{0})) = 0 < 1$ at the zero-effort profile. By their criterion, *zero effort is Pareto efficient in a pure forest domain*: there is no mutual gain from any pattern of effort that can propagate through the network.

This seems paradoxical—if there are positive spillovers, how can doing nothing be efficient? The resolution is that on a DAG, no effort ever "completes a cycle." Spillovers flow in one direction only; no action can benefit everyone along a feedback loop. Without cycles, the collaborative surplus that justifies collective action is absent.

The design problem escapes this trap precisely by creating internal structure within tiles. Inside a tile, all nodes share a single regulator; the within-tile game is cooperative by assumption, not a strategic interaction. A coarser tiling creates larger "cooperative units" while keeping the inter-tile links acyclic. The welfare gain from coarsening under $\alpha < 0$ (the Oates regime) is not about generating EG-style cooperative surplus—it is about eliminating the cross-tile harm that forces downstream over-provision. The relevant criterion remains the sign of α relative to α^*, not the spectral radius. But the EG observation explains why fine tilings do not generate cooperative gains of their own: the tile-level game is always a DAG, and DAGs cannot sustain the cyclic feedback that makes collective action Pareto-improving. The design question is which DAG to play on—not whether to cooperate, but how to partition the lack of cooperation most efficiently.

The Boolean structure and what it bought. The reader who has followed the proofs will have noticed that the lattice-theoretic language—Boolean algebras, Galois connections, closure operators—was not decorative. It did three pieces of work that a less structured formulation could not have done as cleanly.

First, separability. The design problem over $\mathcal{W}(V, \to)$ is a problem over 2^G: a binary choice at each non-root node $v \in G$. The welfare gradient $b(\bar{\gamma}) - \Delta G_v$ at each node depends only on the local topology of the tree (the depth of v, the provision at the root tile downstream of v), not on the choices made at other nodes. This separability—the absence of cross-boundary spillovers in the *design* problem, as opposed to the *provision* problem—is a direct consequence of $\mathcal{W} \cong 2^G$. In a non-Boolean tiling lattice (e.g., the partition lattice of a non-forest domain, where a tile can have multiple outlets), boundary activations interact: opening boundary v changes the quotient domain in ways that affect the welfare gradient at boundary w. The Boolean structure eliminates these interactions and reduces the design problem to $|G|$ independent one-dimensional choices.

Second, the Galois connection. The anti-tone Galois connection between $(\mathcal{W}, \preceq)$ and $(M, \leq)$ exists precisely because the boundary activation decisions are independent. The proof of Proposition 3.11 was four lines; the substance is entirely in the Boolean structure. For a non-forest domain—a river delta with multiple channels, or a road network where several routes connect the same nodes—the watershed condition fails, the partition lattice is no longer Boolean, and the duality breaks. There is no clean per-boundary break-even grant and no Galois connection of this form. The fiscal federalism literature has long asked whether map design and grant design are substitutes or complements; the present answer is: they are exactly dual when the domain is a forest, and the notion of exact duality breaks when it is not.

Galeotti, Golub, and Goyal (2020) study a complementary problem: given a fixed game graph, how should a planner allocate a budget of incentive transfers to maximize welfare? They show that for games of strategic substitutes ($\alpha > 0$ in our notation), the optimal large-budget intervention concentrates on the bottom eigenvector of the adjacency matrix. In a directed forest, the bottom eigenvector in the undirected sense has its largest components at the leaves—the most upstream nodes. GGG's prescription is thus: target the most upstream jurisdictions. This is independently the conclusion of the break-even grant analysis in §3.4, where $\mu_v^* = |\alpha| - \alpha^*$ is the same for all boundaries v in the two-tile chain but in general grows with the distance of v from the root (deeper boundaries have more upstream provision flowing past them, raising the break-even threshold). The two approaches are complementary: GGG optimizes targeting on a fixed tiling; the Galois connection optimizes the tiling itself. Together they provide a complete characterization of the planner's two instruments and their interaction—the Galois connection says when to redesign the map, and GGG says how to target grants on whatever map remains.

Third, the envelope construction. The lowest common outlet $d(Q)$ used in §3.5 is well-defined because every non-empty subset of a tree has a unique most-downstream element. In a directed acyclic graph with multiple sinks, or in a lattice with forks, this fails: there may be no unique $d(Q)$ and the watershed envelope is not a single tile. Again, the forest condition is not just sufficient but necessary for the construction to be canonical.

Taken together, these observations support a stronger claim than "the forest domain is a convenient special case." The forest domain is the *maximal* class on which the partition theory of §2, the welfare monotonicity of §3.3, the Galois duality of §3.4, and the envelope projection of §3.5 all simultaneously hold. Every river network is a forest domain; every forest domain can be interpreted as a river network. The theory is not general with some applications: it is built for this class of objects, and the class is exactly right.

4 Extensions

Section 4.1 closes the gap in Proposition 3.15: it defines welfare for arbitrary partitions and completes the proof that watershed-aligned tilings dominate. Section 4.2 takes a different direction, dropping the transitivity assumption and connecting the model to the local-spillover literature.

4.1 Games on Non-Watershed Partitions

Proposition 3.15 asserts $W(\mathcal{P}^{ws}) \geq W(\mathcal{P})$ but relies on $W(\mathcal{P})$ being defined for an arbitrary partition $\mathcal{P}$. For $\mathcal{P} \in \mathcal{W}(V, \to)$ the definition is §3.2; for non-watershed $\mathcal{P}$ it is not. This section supplies the missing definition and completes the proof.

The single-outlet convention. Fix a linear extension ℓ of the partial order $(V, \to)$—a total order on V such that $v \to w$ implies $\ell(v) < \ell(w)$, so downstream nodes receive smaller labels. For each tile T in any partition $\mathcal{P}$ of V, define the *designated outlet*

$$d_T \;:=\; \arg\min_{t \in T} \ell(t),$$

the most-downstream element of T under ℓ. For watershed tiles, d_T equals the unique outlet from Definition 2.2. For non-watershed tiles, the linear extension breaks any ambiguity in the partial order.

Definition 4.1 (Extended provision game)**.** Let $\mathcal{P}$ be any partition of V and fix a linear extension ℓ. The *extended provision game* $\Gamma(\mathcal{P}, \alpha)$ is the game of Definition 3.1 played on $\mathcal{P}$ with designated outlets $\{d_T\}_{T \in \mathcal{P}}$: the inter-tile relation is $T_i \to_{\mathcal{P}} T_j$ iff $d_{T_i} \to d_{T_j}$ in $(V, \to)$ and $T_i \neq T_j$; the composite at T_j is $\gamma_{T_j} = g_{T_j} + \alpha \sum_{T_i \to_{\mathcal{P}} T_j} g_{T_i}$; the payoff is $b(\gamma_{T_j}) - g_{T_j}$. Equilibrium welfare $W(\mathcal{P})$ is the utilitarian sum at the unique Nash equilibrium.

A key structural observation: the inter-tile graph of any partition under the single-outlet convention is a forest. Since outlets d_T are nodes of the forest $(V, \to)$, the relation $d_{T_i} \to d_{T_j}$ inherits the forest property—each tile has at most one direct downstream neighbor. Proposition 3.2 (existence and uniqueness) and Lemma 3.5 (constant composite $\gamma^*_T = \bar{\gamma}$) therefore hold for $\Gamma(\mathcal{P}, \alpha)$ for *any* partition $\mathcal{P}$. The entire welfare analysis of §3—including the covering-step formula—applies to the extended game.

Completing the proof of Proposition 3.15. The key structural observation is that nodes belonging to the *non-outlet component* of a cross-tree tile—those in a tree whose root is not the tile's designated outlet—are invisible to the provision game in $\mathcal{P}$: because the tile's outlet lies in a different tree, the single-outlet convention produces no inter-tile edge connecting those nodes' tree to anything downstream, so their provision contributes only to their own tile's composite. After tree-alignment they become part of a within-tree tile; after watershed-coarsening they are absorbed into the upstream cone of their local outlet. The game structure in each tree is unchanged by this process, so the welfare contribution of non-outlet components is the same in $\mathcal{P}$ and $\mathcal{P}^{ws}$.

What *does* change is the treatment of the *within-tree non-watershed* inter-tile edges of $\mathcal{P}$: edges $Q \to_{\mathcal{P}} T_j$ where d_Q and d_{T_j} lie in the same tree but $Q \neq \mathcal{T}(d_Q)$. Let E_{orig} denote this set. Watershed-coarsening (Step 2) merges each such edge away. Each merger is a within-tree covering step to which the formula $W(\text{merged}) - W(\text{split}) = \Delta G - b(\bar{\gamma})$ applies. For $\alpha < 0$ and $|\alpha| > \alpha^*$, splitting increases the over-provision burden on the downstream root tile (the same mechanism as equation (2)), giving $\Delta G > b(\bar{\gamma})$ and hence a positive welfare gain from each merger.

Proof of Proposition 3.15 (completed). Under the single-outlet convention, $W(\mathcal{P})$ is well-defined and the inter-tile graph is a forest. Let E_{orig} be the within-tree non-watershed inter-tile edges of $\mathcal{P}$.

Non-outlet components are welfare-neutral. For each cross-tree tile $T \in \mathcal{P}$ with outlet d_T in tree τ, the nodes of T lying in any other tree $\tau' \neq \tau$ generate no inter-tile edges in $\mathcal{P}$ (since $d_T \notin \tau'$, no edge $d_T \to d_{T_j}$ can be within τ'). Tree-alignment splits T at tree boundaries; watershed-coarsening merges the resulting τ'-pieces into their upstream cones. By the argument of the previous paragraph, these operations leave the game structure in every tree unchanged, contributing zero net welfare change.

Within-tree merges raise welfare. Each edge $e \in E_{\text{orig}}$ is eliminated by a covering merge in Step 2. In the within-tree inter-tile chain of $n \geq 2$ tiles, Corollary 3.6 gives equilibrium provision $\bar{\gamma}(1 + |\alpha| + \cdots + |\alpha|^j)$ at the tile j steps from the root; the top tile ($j = n-1$) provides $\bar{\gamma}(1 + |\alpha| + \cdots + |\alpha|^{n-1})$, and a covering merge absorbs it, so $\Delta G = \bar{\gamma}(1 + |\alpha| + \cdots + |\alpha|^{n-1})$. Since $n \geq 2$, this gives $\Delta G \geq \bar{\gamma}(1 + |\alpha|) > \bar{\gamma}(1 + \alpha^*) = b(\bar{\gamma})$ and hence welfare gain $\Delta G - b(\bar{\gamma}) > 0$ whenever $|\alpha| > \alpha^*$.

The construction $\mathcal{P} \to \mathcal{P}^{ws}$ proceeds through a finite sequence of covering merges, each raising welfare. By telescoping, $W(\mathcal{P}^{ws}) \geq W(\mathcal{P})$, with equality iff every step is degenerate—i.e., $E_{\text{orig}} = \emptyset$. $\square$

Remark 4.2 (Why the two-step decomposition fails)**.** The natural decomposition—tree-alignment raises welfare, then within-tree coarsening raises welfare—is wrong about the first step. Tree-alignment splits cross-tree tiles into per-tree pieces, creating new within-tree inter-tile edges. For $\alpha < 0$ with $|\alpha| > \alpha^*$, each new edge forces the downstream tile to over-provide, reducing welfare; Step 2 then merges those same edges away, recovering the loss. The net effect of creating and then removing a new edge is zero: only the original within-tree non-watershed edges E_{orig} of $\mathcal{P}$ contribute a net gain, each one positive when $|\alpha| > \alpha^*$.

4.2 Non-Transitive Domains

The forest domain $(V, \rightarrow)$ requires transitivity: if $v \rightarrow w$ and $w \rightarrow x$ then $v \rightarrow x$, so provision at v reaches every downstream node in the chain. This models cumulative spillovers well (a pollutant released upstream accumulates downstream), but poorly for networks where spillovers attenuate at each link and do not compound. This section relaxes transitivity and studies the resulting local-spillover game.

Definition 4.3 (Local-spillover domain)**.** A *local-spillover domain* is a pair $(V, \rightharpoonup)$ where V is a finite set and $\rightharpoonup$ is an irreflexive, antisymmetric relation on V—a directed acyclic graph (DAG) without self-loops. We read $v \rightharpoonup w$ as "v directly spills into w." Transitivity is not required: $v \rightharpoonup w$ and $w \rightharpoonup x$ need not imply $v \rightharpoonup x$.

In the local-spillover game on $(V, \rightharpoonup)$ with a watershed tiling $\mathcal{T} \in \mathcal{W}(V, \rightarrow^+)$ (where $\rightarrow^+$ is the transitive closure of $\rightharpoonup$), the composite at tile T reflects only direct upstream neighbors:

$$\gamma_T \;=\; g_T + \alpha \sum_{\substack{T' \in \mathcal{T}: \\ d_{T'} \rightharpoonup d_T}} g_{T'}.$$

The Constant Composite Lemma (Lemma 3.5) holds tile-by-tile (each tile's first-order condition is $b'(\gamma_T^*) = 1$ regardless of how the composite is assembled), so $\gamma_T^* = \bar{\gamma}$ for all T. What changes is the equilibrium provision vector $\mathbf{g}^*$: tiles deep in the chain respond only to their direct upstream neighbor, not the full upstream chain, so they free-ride less and provide more.

Observation 4.4. If $(V, \rightharpoonup)$ is a local-spillover domain whose underlying DAG is a forest (each node has at most one direct downstream neighbor under $\rightharpoonup$), then the transitive closure $(V, \to^{+})$ is a forest domain in the sense of Section 2. Hence $\mathcal{W}(V, \to^{+}) \cong B_{|V|-k}$, and all of §§2–4.1 applies to the tiling problem on $(V, \to^{+})$.

The local-spillover game on $(V, \rightharpoonup)$ and the transitive game on $(V, \to^{+})$ share the same watershed lattice as their design space but differ in equilibrium welfare for intermediate tilings.

The poles are unchanged. At $\hat{0}$ (single tile), the planner internalizes all spillovers regardless of whether they are transitive or local: the first-best provision profile is feasible in both games. At $\hat{1}$ (all singletons), the duplication structure depends on the local graph: tiles that are not directly linked in $\rightharpoonup$ face no composite from each other, so their provisions are determined independently. In a long chain $v_1 \rightharpoonup v_2 \rightharpoonup \cdots \rightharpoonup v_n$, the local game at $\hat{1}$ has each v_i responding only to v_{i-1}, shortening the effective strategic cascade. Total provision under local spillovers at $\hat{1}$ is *higher* than under transitive spillovers (each tile free-rides less, since it receives a smaller composite from upstream). Welfare at $\hat{1}$ is correspondingly lower.

Monotonicity fails at intermediate tilings for $\alpha < 0$. For transitive spillovers, every covering step in $\mathcal{W}(V, \to^{+})$ changes welfare in the same direction (Proposition 3.7 and the envelope analysis of §3.5): $\operatorname{sign}(\alpha)$ determines whether finer or coarser is uniformly better. For local spillovers this uniformity breaks. The welfare gain from opening a boundary at v depends on how many direct-link hops separate v's tile from the nearest upstream provision. Each additional hop attenuates the internalization benefit by one factor of $|\alpha|$. The result is a family of per-boundary thresholds that grow with depth in the chain: boundaries deeper in the local-spillover graph require a larger $|\alpha|$ to justify. When $\alpha < 0$ and $|\alpha|$ lies in the intermediate range $|\alpha| < \alpha^{*} < |\alpha|(1 + |\alpha|)$, the first boundary opened improves welfare but a second does not, violating monotonicity.

Lemma 4.5 (Monotonicity failure under local spillovers)**.** *There exists a local-spillover domain $(V, \rightharpoonup)$ with forest-DAG structure, spillover parameter $\alpha < 0$, and bliss-point surplus ratio α^{*} satisfying $|\alpha| < \alpha^{*} < |\alpha|(1 + |\alpha|)$, for which the welfare function W_{local} is not monotone on $\mathcal{W}(V, \to^{+})$: there exist tilings $\mathcal{T} \prec \mathcal{T}'$ ($\mathcal{T}'$ finer) with $W_{local}(\mathcal{T}') < W_{local}(\mathcal{T})$.*

Proof. Take the four-node chain $a \rightharpoonup b \rightharpoonup c \rightharpoonup d$ (root d; $G = \{a, b, c\}$; transitive closure $(V, \to^+)$ is the chain $a \to b \to c \to d$). Consider three tilings: $\hat{0} = \{a, b, c, d\}$, $\mathcal{T}_1 = \{a\}, \{b, c, d\}$, and $\mathcal{T}_2 = \{a\}, \{b, c\}, \{d\}$, with lattice ordering $\hat{0} \prec \mathcal{T}_1 \prec \mathcal{T}_2$ (each step finer by one boundary).

In the local game, only direct links carry spillover. The inter-tile composite at each tile is determined by its unique direct upstream neighbor (if any).

Welfare at $\hat{0}$: single tile, $g = \bar{\gamma}$, $W(\hat{0}) = b(\bar{\gamma}) - \bar{\gamma} = \bar{\gamma}\alpha^*$.

Welfare at $\mathcal{T}_1$: tile $\{b, c, d\}$ has outlet d; tile $\{a\}$ links directly to it ($a \rightharpoonup b$, $b \in \{b, c, d\}$). Setting $g_a = \bar{\gamma}$ and $g_{bcd} = \bar{\gamma}(1 - \alpha) = \bar{\gamma}(1 + |\alpha|)$,

$$W(\mathcal{T}_1) = 2b(\bar{\gamma}) - \bar{\gamma}(2 - \alpha) = \bar{\gamma}(2\alpha^* - |\alpha|).$$

Since $|\alpha| < \alpha^*$, $W(\mathcal{T}_1) > \bar{\gamma}\alpha^* = W(\hat{0})$.

Welfare at $\mathcal{T}_2$: tile $\{a\}$ links to $\{b, c\}$ ($a \rightharpoonup b$); tile $\{b, c\}$ links to $\{d\}$ ($c \rightharpoonup d$); tile $\{a\}$ does *not* link to $\{d\}$ (no direct $a \rightharpoonup d$). Setting composites equal to $\bar{\gamma}$ gives $g_a = \bar{\gamma}$, $g_{bc} = \bar{\gamma}(1 + |\alpha|)$, $g_d = \bar{\gamma}(1 + |\alpha| + |\alpha|^2)$, and

$$W(\mathcal{T}_2) = 3b(\bar{\gamma}) - \bar{\gamma}(3 + 2|\alpha| + |\alpha|^2) = \bar{\gamma}(3\alpha^* - 2|\alpha| - |\alpha|^2).$$

The welfare change from $\mathcal{T}_1$ to $\mathcal{T}_2$ is

$$W(\mathcal{T}_2) - W(\mathcal{T}_1) = \bar{\gamma}\big(\alpha^* - |\alpha|(1 + |\alpha|)\big) < 0,$$

since $\alpha^* < |\alpha|(1 + |\alpha|)$ by assumption. Thus $\mathcal{T}_1 \prec \mathcal{T}_2$ (finer) but $W_{local}(\mathcal{T}_2) < W_{local}(\mathcal{T}_1)$, violating monotonicity. ☐

Remark 4.6 (Depth-dependent thresholds)*.* The proof identifies the mechanism precisely. The covering step $\hat{0} \to \mathcal{T}_1$ has welfare gain $\bar{\gamma}(\alpha^* - |\alpha|)$—the same formula as in the transitive game, because the two-tile inter-tile chain has a single direct link. The next step $\mathcal{T}_1 \to \mathcal{T}_2$ has gain $\bar{\gamma}(\alpha^* - |\alpha|(1 + |\alpha|))$: the threshold has risen by a factor of $1 + |\alpha|$ because the new downstream tile $\{d\}$ receives only from $\{b, c\}$ (one hop), not from $\{a\}$ (which reaches $\{d\}$ only transitively). In the transitive game both tiles' provisions contribute; in the local game only the direct upstream tile does, attenuating the internalization benefit. The k-th covering step in a chain has threshold $|\alpha|(1 + |\alpha| + \cdots + |\alpha|^{k-1}) = |\alpha|(1 - |\alpha|^k)/(1 - |\alpha|)$, which grows with depth k. The design rule for local spillovers is therefore boundary-by-boundary: activate boundary v if and only if $|\alpha|$ exceeds the threshold for v's depth in the local-spillover chain, not simply whether $|\alpha| > \alpha^*$.

Connection to the local public goods literature. Bramoullé and Kranton (2007) study exactly the local-spillover game on undirected networks with binary choices: their Nash equilibria are maximal independent sets, and the optimal network concentrates provision on "specialists" who are maximally exposed to each other. The smooth-payoff directed analogue here replaces maximal independent sets with the depth-dependent covering-step calculation: the "specialists" are tiles whose outlets are directly linked in $\rightharpoonup$. Lemma 4.5 and the accompanying Remark are the smooth analogue of the Bramoullé–Kranton insight: optimal boundary placement tracks the actual spillover graph $\rightharpoonup$, not its transitive closure $\to^+$, and the welfare cost of a misplaced boundary grows with its depth in the local-spillover chain. Connecting these depth-dependent thresholds to the Galois connection and grant duality of §3.4 is left for future work.

5 Conclusion

This paper developed a lattice-theoretic framework for jurisdiction design under ordered spillovers. The watershed lattice $\mathcal{W}(V, \to) \cong B_{|V|-k}$ captures all admissible maps as independent binary switches; the provision game's equilibrium welfare is monotone on this lattice with direction determined by the sign of α; and the Galois connection between maps and grant schedules gives every off-lattice partition a precise fiscal comparator. The three results rest on a single structural constant:

$$\alpha^* = \frac{b(\bar{\gamma}) - \bar{\gamma}}{\bar{\gamma}},$$

the bliss-point surplus ratio. It is the monotonicity threshold, the break-even grant per boundary, and the watershed-dominance condition. We close with the open questions it raises.

Calibration. Whether any specific jurisdictional border is in the Oates regime ($|\alpha| > \alpha^*$) is an empirical question: does the spillover transfer rate across that border exceed the bliss-point surplus ratio? Both quantities are in principle observable—$\bar{\gamma}$ from equilibrium provision data, $b(\bar{\gamma}) - \bar{\gamma}$ from cost-benefit analysis of provision programs, and $|\alpha|$ from downstream monitoring of transfer rates. Calibrating α^* for specific networks—Chesapeake Bay nutrient loads, Rhine basin water quality, Great Lakes pollution standards—would identify which existing borders impose costs above the break-even threshold and quantify the fiscal gap $\sum_{v\in\phi(\mathcal{P}^{ws})\setminus\phi(\mathcal{P})} \mu_v^*$.

Edge-indexed spillovers. The homogeneous parameter α assigns the same transfer rate to every inter-tile boundary. Real spillover networks are heterogeneous: a pollutant attenuates differently on different river reaches, and point and diffuse sources have distinct profiles. With edge-indexed spillovers $\{\alpha_e\}_{e\in E(\mathcal{T})}$, each boundary v has its own break-even threshold $\alpha_e^* = (b(\bar{\gamma}) - \bar{\gamma})/\bar{\gamma}$—the same formula, but applied separately. The design problem becomes boundary-by-boundary: activate the split at v iff $|\alpha_e| > \alpha_e^*$. This is the empirically tractable version of the design problem, and it is exactly what the Boolean structure of $\mathcal{W}$ prescribes: each of the $|V| - k$ switches is decided independently. Galeotti, Golub, and Goyal (2020) provide the spectral analog for continuous interventions; connecting their eigenvector targeting to the boundary-by-boundary threshold here is an open problem.

The political economy premium. Besley and Coate (2003) show that centralized provision under legislative bargaining adds a political economy cost above the Pigouvian optimum. In our framework this raises the effective threshold: a border requires $|\alpha| > \alpha^* + \delta_{BC}$ for some political economy premium $\delta_{BC} > 0$ before centralization improves welfare. Measuring δ_{BC} for real jurisdictions—the gap between the economic threshold α^* and the transfer rate at which jurisdictions empirically centralize— gives a structural estimate of the political economy cost of bad borders, distinct from the direct Oates welfare gap.

Dynamics. The model is static. In a repeated game, each period's provision creates a stock of environmental quality; the discount rate and detection probability jointly shift the effective threshold. The Chesapeake Bay Program's 40-year trajectory suggests that the dynamic threshold matters: borders that appear in the Oates regime under static analysis may not be under discounted dynamic welfare once transition costs and irreversibilities are included. Whether $\mathcal{W}(V, \to)$ has a natural dynamic analog—a Markov chain on the Boolean algebra, with covering steps as transitions and welfare as a Lyapunov function—is an open structural question.

Fork domains. Every non-watershed tile Q in a *forest* domain is *decomposable*: Q can be partitioned into watershed sub-tiles by splitting along branch boundaries. The single-outlet convention of §4.1 exploits this—it is a clean approximation precisely because decomposability holds. Decomposability *fails* in fork domains, where a node v satisfies $v \to w$ and $v \to x$ for distinct w, x. Fork node v belongs to both downstream sub-basins; assigning it to one branch severs its spillover to the other. Real river deltas and distributary networks have this structure, and extending the model to fork domains—replacing the Boolean algebra $\mathcal{W}$ with a product-lattice that tracks sub-basin allocation within each fork—is a natural next step.

References

V. Bala and S. Goyal. A noncooperative model of network formation. *Econometrica*, 68:1181–1229, 2000.

O. Baetz. Social activity and network formation. *Theoretical Economics*, 10:315–340, 2015.

M. Belhaj, S. Bervoets, and F. Deroïan. Efficient networks in games with local complementarities. *Theoretical Economics*, 11:357–380, 2016.

T. Bergstrom, L. Blume, and H. Varian. On the private provision of public goods. *Journal of Public Economics*, 29:25–49, 1986.

T. Besley and S. Coate. Centralized versus decentralized provision of local public goods: a political economy approach. *Journal of Public Economics*, 87:2611–2637, 2003.

F. Bloch and U. Zenginobuz. The effect of spillovers on the provision of local public goods. *Review of Economic Design*, 11:199–216, 2007.

Y. Bramoullé and R. Kranton. Public goods in networks. *Journal of Economic Theory*, 135:478–494, 2007.

R. J. Carroll, D. M. Konisky, and C. Reenock. Policy Devolution and Cooperation Dilemmas. Working paper, 2025.

B. A. Davey and H. A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, 2nd edition, 2002.

M. Elliott and B. Golub. A network approach to public goods. *Journal of Political Economy*, 127:730–776, 2019.

B. Fong and D. I. Spivak. *An Invitation to Applied Category Theory*. Cambridge University Press, 2019.

A. Galeotti, B. Golub, and S. Goyal. Targeting interventions in networks. *Econometrica*, 88:2445–2471, 2020.

T. Hiller. Peer effects in endogenous networks. *Games and Economic Behavior*, 105:349–367, 2017.

M. O. Jackson and A. Wolinsky. A strategic model of social and economic networks. *Journal of Economic Theory*, 71:44–74, 1996.

X. Li. Designing weighted and directed networks with complementarities. *Games and Economic Behavior*, 140:556–574, 2023.

W. E. Oates. *Fiscal Federalism.* Harcourt Brace Jovanovich, New York, 1972.

W. E. Oates. An essay on fiscal federalism. *Journal of Economic Literature*, 37:1120–1149, 1999.

Y. Sun, W. Zhao, and J. Zhou. Sequential network design. *Journal of Economic Theory*, 232:106145, 2026.

C. M. Tiebout. A pure theory of local expenditures. *Journal of Political Economy*, 64:416–424, 1956.