

The Hydropower Promise? Local Governance and Rural Redistribution

Nicole Baerg

April 2026

Abstract

Representatives controlling large discretionary funds are expected to direct resources towards their own local communities, and access to resource windfalls is expected to intensify this behaviour. The existing literature finds that external accountability mechanisms, such as elections, media scrutiny, and rule-based distribution, can offset some of this. We argue that internal accountability mechanisms are equally if not more valuable yet underexplored. When board seats are allocated one per community across communities competing for the same fixed budget, home bias becomes costly because the coalition support needed to extract resources homeward does not emerge. This works whether officials are elected or appointed, but only when board members face zero-sum decisions. We test this theory using the Columbia Basin Trust, a hydropower benefit-sharing fund in British Columbia. We combine the synthetic control method with random-effects models of grant allocation and seven biennial waves of resident survey data. We find positive developmental returns and no evidence of home bias. We conclude that mixed-territorial seat allocation is a low-cost institutional design choice with measurable consequences for the distribution of windfall revenues, while flagging the limits of what the data can establish about director-level mechanisms.¹

¹Thank you to Sergio Ascencio, John Janmaat, Jamye Jones, Nikolay Marinov, Laura Montanaro, Johanna Rickne, Jake Ricks, Christopher Schulz, Jon Slapin, Martin Steinwand, Joanne Taylor and John Wagner for comments and questions. The author used Claude (Anthropic) to assist with drafting and editing, and to suggest candidate references for verification. All references and factual claims were independently verified by the author, who is responsible for the content of the paper. This is a working draft. Please do not cite without permission.

1 Introduction

Large energy projects, such as hydropower, wind, and nuclear power, are now widely expected to *deliver* and *develop* the communities they affect. In the United Kingdom, the development of Sizewell C, a nuclear power station, comes with a £23 million community fund targeting East Suffolk, a county with high levels of rural deprivation. In Sweden, the government has allocated over SEK 1 billion to predominantly rural municipalities that approve onshore wind projects to compensate for the burdens of hosting energy infrastructure. Paying communities to accept energy infrastructure does not resolve how those funds are governed, nor does it guarantee the delivery and development outcomes aspired to. In recent decades, governments and international organisations have promoted *benefit-sharing mechanisms* as a way to transform disruptive infrastructure into engines of regional development (Schulz and Skinner, 2022). These arrangements dedicate non-tax, windfall revenues from resource extraction to affected regions and often embed them in participatory local governance institutions. Because these revenues do not arise from taxation, they come with what Ahmed (2012) calls the perils of unearned income, a weakened fiscal contract that would otherwise give citizens an incentive to monitor how money is spent.

Outcomes from these funds vary widely. Alaska’s Permanent Fund Dividend distributes oil revenues as an equal cash payment to every resident. It has no board, no discretion, and therefore no mechanism for targeted local development. The fund achieves equal distribution but cannot convert resource windfalls into productive local community development on its own, and recent research finds limited developmental returns (Goldsmith, 2010), though Loeffler (2023) shows the dividend has increased voter turnout. Norway, alternatively, has both delivered resource funds to the public and developed economically. Mayors distribute hydropower revenues through discretionary local spending, and Andersen and Sørensen (2022) find they do so broadly and without personal political gains. The authors attribute this to electoral accountability supported by high voter turnout and active local media. In Canada, at Muskrat Falls in Labrador, the outcome is again different. The Nunatsiavut Inuit and Innu Nation, Indigenous peoples whose traditional territories cover the Labrador-Quebec peninsula, bore the environmental and social costs of hydropower development while being systematically excluded from the boards that controlled the benefits (Samson, 2016; Yellowhead Institute, 2021).

What kind of governance structure achieves equitable distribution and economic development in the context of energy windfalls? Are external mechanisms of enforcement, such as elections and media scrutiny, necessary, or can other internal accountability mechanisms also protect resources from political or elite capture? We argue that external accountability is not the only answer. When board seats are allocated one per community across communities competing for the same fixed budget, what we call *mixed-territorial seat allocation* (MTSA), the structure of the board itself constrains the tendency of representatives to direct resources toward their own communities. Two conditions are required: board members must face genuinely zero-sum decisions among rival claimants, and the board must combine district-facing members with broader appointees who have no territorial stake.

To intuitively understand MTSA consider the Federal Open Market Committee, the body that conducts U.S. monetary policy. The FOMC has twelve voting members: the seven Governors of the Federal Reserve Board, who are appointed by the President and confirmed by the Senate; the president of the Federal Reserve Bank of New York, who is a permanent voting member; and four of the eleven other regional Reserve Bank presidents, who rotate annually. The Governors have no specific district affiliation; the Reserve Bank presidents do. These different appointment types make the Federal Reserve System a mixed-territorial body. Unlike MTSA systems, however, members on the FOMC do not face zero-sum decisions in competition with other Bank Presidents. The FOMC decides on monetary policy, not district-specific budgets, and the 5th District, Richmond, does not lose when the 12th District, San Francisco, gains because there is no fixed pot that district members compete over. In this paper, we consider MTSA institutions that have zero-sum rivalry. The same communities that nominate regional directors also compete for a fixed grant budget such that every dollar directed to district A is a dollar not going to district B.

Any director who wants to steer grants homeward needs a majority. In the legislative case, research shows that to build this majority she would build a coalition of others also wanting to raise funds, with the consequence that the pot grows with it (i.e. pork barrel politics). In MTSA with rivalry, while other regional directors also have an incentive to redirect funds to their own community, the board is also comprised of central appointees with no district stake and therefore no reason to join any regional extraction scheme. If the pot is fixed, members cannot accommodate the demands of co-opted members. As a consequence of the mixed appointment types and the zero-sum nature of resource-fund distribution, a majority in favour of extraction cannot form. Regional members face structural constraints in directing funds homeward, and political capture is constrained.

We test this theory using the Columbia Basin Trust (CBT), a benefit-sharing fund established in British Columbia, Canada, in 1995 to compensate communities for the costs of reservoir construction under the Columbia River Treaty. The CBT board has twelve members. Five seats are nominated by regional districts and one is nominated by the Ktunaxa Nation, a First Nation whose traditional territory covers part of the Columbia Basin. The remaining six seats are appointed directly by the Province with no district stake but with a Basin-wide interest. All twelve members make annual budget allocation decisions. The board is mixed-territorial: it combines regional nominees competing to bring funds home with provincial appointees who answer to no specific community.

Empirically, we use three strategies. First, the synthetic control method estimates the CBT's aggregate effect on regional economic development, operationalised as new business registrations per development region. Second, a random-effects model of grant allocation tests whether districts represented by politically experienced or long-tenured directors receive disproportionate funding. Third, seven biennial waves of resident survey data (2010–2022) examine whether residents perceive the allocation as fair across the Basin.

We find a small but positive cumulative effect on regional economic development, amounting to approximately 560 additional business registrations between 1995 and 2019, equivalent to one additional year of growth over two decades. At the district level, we find no evidence of home bias. Po-

litically experienced and long-tenured directors are associated with lower, not higher, home-district grant receipts, though within-district variation in political experience is thin and the coefficients on these terms are not statistically distinguishable from zero. Resident surveys show no perceptions of systematic regional disadvantage across the Basin: across all four quadrants and all seven waves, residents rate the CBT's contribution to their community at 70% positive or above.

Taken together, the results suggest that external enforcement mechanisms such as elections and media scrutiny are not the only way to protect resource windfall revenues from political capture. Internal governance structures matter, and mixed-territorial seat allocation is one understudied and promising mechanism. As shown, when every regional director is also a rival, no director can build the majority needed to steer grants homeward, and provincial appointees with no district stake have no reason to join an extraction scheme. We are careful, however, to note what the data cannot establish. The grant-allocation result is suggestive given the small panel (six districts, twenty-two years) and the limited within-district variation in director political experience. The strongest evidence is the descriptive one: no district dominates grant receipts in any consistent way, and no quadrant of the Basin perceives itself as systematically disadvantaged. For the growing class of appointed benefit-sharing bodies being created alongside renewable energy development worldwide, MTSA is a low-cost design choice that warrants serious attention.

2 Board Design as an Alternative to External Accountability

The Resource Curse at the Sub-national Level

The development problem. The resource curse literature documents a consistent pattern. Windfall revenues from resource extraction are associated with rent-seeking, misallocation, and weak developmental returns at the national level (Sachs and Warner, 1995; Havranek et al., 2016) and increasingly at the sub-national level as well (Brollo et al., 2013; Nikolova and Marinov, 2017). The mechanism runs through the contract between citizens and government (Ross, 2001). When citizens pay taxes, they have a direct incentive to monitor how that money is spent and to punish officials who spend it badly. This monitoring incentive is the foundation of fiscal accountability. Windfall revenues, including remittances, government transfers, natural-resource revenues, and disaster funds, break this link. When governments receive money from oil, hydropower, or other resource rents rather than from citizens directly, they have less reason to justify spending decisions. Citizens have fewer grounds to demand accountability. The result is that windfall revenues are systematically more likely to be captured, wasted, or misallocated than equivalent revenues raised through taxation.

The problem is more severe at the sub-national level. Sub-national governments tend to operate with weaker accountability institutions: fewer independent auditors, less media scrutiny, lower electoral turnover, and lower capacity for organised citizen monitoring. The windfall itself can make matters worse. Brollo et al. (2013) show that larger intergovernmental transfers increase corruption and reduce the quality of elected officials precisely because windfalls reduce voter incentives to

monitor incumbents. Nikolova and Marinov (2017) document this directly, showing that Bulgarian municipal politicians systematically misappropriate windfall disaster relief funds, with the largest effects in communities with weaker civic capacity. Havranek et al. (2016) find that the negative developmental effects of resource windfalls are significantly larger where institutional quality is weak, a condition that holds even in wealthy advanced economies, as shown by the governance histories of the Alaska Permanent Fund and Muskrat Falls. Distributional distortions follow a similar logic: Fournaies and Mutlu-Eren (2015) show that central governments in England allocate significantly more to locally co-partisan councils, and the channel running through partisan electoral incentives is the same channel that exists in the resource-curse literature. The discipline that voters and partisan competition impose on incumbents is external to the governance body itself. When that external infrastructure is weak or absent, the disciplining force is weak or absent with it, and remedies that depend on it are correspondingly costly to implement.

The distributional problem. Even when windfall revenues are invested rather than consumed, funds tend to flow to politically connected places rather than being distributed equally. When funds are specifically directed to the territory the representative is from, the literature refers to this as home bias. Fenno (1978) shows that representatives systematically cultivate relationships with their home constituencies and claim credit for grants directed there. Mayhew (1974) argues that the incentive structure of representative democracy makes this nearly inevitable: legislators engage in credit-claiming by linking public expenditures to their own actions, building local visibility that sustains electoral careers. Ferejohn (1974) and Shepsle and Weingast (1981) formalise the institutional conditions under which distributive coalitions form and sustain themselves.

Home bias need not disappear under appointments rather than elections. Appointments remove the electoral incentive to claim credit, but reputational concerns remain. Appointed board members are typically nominated by the local districts or areas they represent, and consequently are embedded in local political and elite networks. While Enikolopov (2014) shows that appointed chief executives in US local governments are less likely than elected ones to engage in targeted redistribution, the difference is one of degree and career incentive rather than kind. Even on the FOMC, researchers show that Bank Presidents speak more about inflation and unemployment when their home districts are suffering those conditions (Baerg, 2020).

What, if any, are plausible solutions? The existing literature focuses on two main mechanisms of accountability: electoral accountability and rules-based distribution. Both are external to the board. In what follows we introduce a theory that shows when internal board-governance mechanisms can also be effective.

Electoral accountability. The threat of elections can discipline politicians, and this is most likely to occur when voters can observe how windfall revenues are spent and punish officials who misallocate them. Andersen and Sørensen (2022) show that Norwegian mayors do not personally benefit from hydropower windfalls and attribute this to high voter turnout and active local newspaper coverage in small municipalities. The key feature of the Norwegian case is that the community receiving the windfall is the same community that elects the mayor controlling discretionary

spending. If the mayor diverts funds toward patronage, voters notice and can respond. The perils of unearned income are solved by putting funds directly into the municipal budget. Once in the budget, voters behave as if windfalls are tax revenues even though they are not.

One issue is that many of these funds redistribute to rural rather than urban areas. Flavin (2012) shows that elected officials are systematically less responsive to poorer constituencies. Rural communities affected by hydropower, wind, and nuclear projects typically cannot deliver decisive votes in state or national elections, which can make them unimportant to politicians. Electoral accountability works when the people who receive the windfall are the same people who elect the officials controlling it. In most energy infrastructure settings, they are not.

Rules-based distribution. An alternative to electoral accountability is to remove discretion entirely. Alaska’s Permanent Fund Dividend distributes oil revenues to residents as an equal per-capita cash payment by statutory formula, with no board and no discretion over allocation. In 2025, the dividend was \$1,000. There is little evidence that this payment supports economic development (Goldsmith, 2010), but research shows the dividend increased voter turnout, ruling out civic disengagement as a cost of formula-based distribution (Loeffler, 2023). Denmark requires by statute that developers offer at least 20% equity in new wind projects to local residents, with benefits flowing as dividends proportional to shares held. Scotland’s framework for onshore wind community benefits follows a similar logic at the project level: developers pay a standard rate of £5,000 per installed megawatt per year, index-linked for the project lifetime. England is now moving toward making equivalent payments mandatory. Norway’s Glomma and Laagen funds achieve equity through statutory taxation on hydropower capacity, removing board discretion entirely (Lillehammer, 2011).

The lesson across these cases is a trade-off: the conditions that prevent political or elite capture also prevent the targeting of funds for development. Development requires the ability to direct investment where it is most productive: to specific communities, sectors, and infrastructure gaps. In hybrid systems, the Scottish wind fund formula sets the total payment but leaves allocation to local boards, which is precisely where discretion re-enters and home bias can return. The Alaska dividend distributes purchasing power to individuals. It does not fund the business loans, community infrastructure, or local investment that converts a windfall into lasting economic growth.

Mixed-territorial seat allocation (MTSA)

Rather than look to external mechanisms of accountability, we argue that internal board-level institutional structures can provide similar constraints against the resource curse and home bias. The specific institutional arrangement we consider is what we term mixed-territorial seat allocation (MTSA) with zero-sum rivalry. The concept draws on the literature on mixed-member representation in legislatures (Shugart and Wattenberg, 2001) and on territorial seat allocation in federal second chambers, but neither literature studies non-elected boards that allocate fixed budgets across the same communities that nominate the directors. MTSA is a board structure in which seats are allocated one per community across the same communities that compete for the board’s discretionary grant budget. Additional appointees, appointed by a higher-level principal with no territorial stake,

participate as well. The key property is not simply that the board contains members from different communities, but that no two community-facing directors share a principal. Each territorial seat represents exactly one community; that community's gain in grants is every other community's loss.

This structure makes extraction coalition formation difficult so long as each district representative has a single seat, the fund size board members must divide is fixed, and extractive coalition formation requires a majority. The logic follows from the legislative bargaining framework of Baron and Ferejohn (1989). In that model, a budget proposer assembles a minimum winning coalition by offering side payments to a subset of members while excluding the rest. The condition enabling extraction is that potential coalition partners can be recruited cheaply and excluded members have a low reservation price. MTSA removes this condition. Assembling a majority requires recruiting directors whose principals (district citizens) would lose from the arrangement. No such coalition is stable because each potential recruit faces the same incentive to defect and redirect funds to their home district instead. The result is a structural stalemate: no director can extract homeward because the votes needed to do so belong to the directors who would lose from doing so. As the board is also made up of alternative appointees with no district stake who have no reason to join any community-extraction scheme, their presence reinforces the stalemate. Like the electoral accountability mechanism, discretion is retained and home bias is structurally constrained. Unlike the external accountability mechanisms, no external structures of accountability are needed.

The mechanism has three scope conditions. First, the community grant budget must be fixed within allocation cycles, and inter-temporal side payments must be difficult to sustain. If the board can expand total spending in a single sitting to accommodate all district demands simultaneously, the zero-sum constraint disappears. If directors can log-roll across budget cycles, supporting each other's districts in alternating years, the inability to hold a coalition goes away. When budgets grow steadily, the temptation to trade favours across time rises.

Second, seats must be allocated one per community with no two directors sharing a principal. If two directors answer to the same district, they can collude, and a partial coalition across two or three districts becomes possible even without no-stake appointee support. The mechanism requires that every regional director be a rival to every other.

Third, community-specific grants must constitute a meaningful share of total spending. If all spending flows through programmes distributed equally across regions by formula, there is no district-targetable budget to compete over and the distributional rivalry the mechanism depends on does not take place.

We test the following hypotheses:

Hypothesis 1 (Development Effect). *In the absence of electoral accountability, unelected boards managing windfall revenues should produce null or negative developmental returns, consistent with the sub-national resource curse. The MTSA mechanism, if operative, should partially offset this, producing a positive aggregate effect on regional economic development.*

Hypothesis 2 (No Home Bias). *If MTSA disciplines behaviour, district-level grant allocations*

should not track the reputational concerns of the directors representing each district. Districts represented by directors with stronger reputational stakes (operationalised below as prior or concurrent electoral experience and longer board tenure) should not receive disproportionate grant receipts.

Hypothesis 3 (Resident Perception). *If MTSA produces equitable allocation in practice, no one should feel worse off.*

The rest of the paper presents empirical evidence in support of or against these hypotheses using the case study of the Columbia Basin Trust.

3 Data and Methodology

Columbia River Treaty and the formation of the CBT

The Columbia River Treaty, signed in 1961 between Canada and the United States, governs flood control and hydroelectricity on the Columbia River and its tributaries. The Columbia Basin is the great watershed of the Canadian Rocky Mountains. The Canadian portion of the Basin alone is roughly the size of Austria, and the entire Basin (Canadian and American sides combined) is roughly the size of Texas. The Canadian portion contributes approximately 40 per cent of average river flows for the entire Basin region. Some of the original treaty conditions expired in 2024 and the treaty has been undergoing renegotiation, though the renegotiation was paused under the second Trump administration. The construction of three large storage reservoirs under the Treaty (Mica, Keenleyside, and Duncan) flooded agricultural land, displaced communities, destroyed First Nations burial grounds, petroglyphs, and pictographs, and eliminated the natural flood regime that had supported rural economies in the Canadian Basin (Penfold, 2013). The benefits of flood control and hydroelectricity flowed primarily to American ratepayers, downstream communities, and BC Hydro, a crown corporation of the provincial government. The Basin communities that bore these costs received nothing in the original agreement.

An earlier attempt at compensation came in 1989, when BC Hydro offered the regional districts grants in lieu of paying regional taxes. The directors rejected the offer, and in 1990 the provincial government and BC Hydro increased and reorganised the grants paid to Basin communities. These pre-1995 grants matter for the empirical analysis: we backdate the synthetic control treatment to 1990 in our robustness checks to test whether they, rather than the 1995 endowment, drive any positive effect.

The CBT itself was created in 1995 through the Columbia Basin Trust Act, after a small group of Basin politicians and community activists organised across party lines and placed allies in the provincial legislature during the 1991 election. The CBT received a \$295 million endowment at its creation, of which \$250 million was committed to regional hydropower projects and \$45 million was earmarked for direct community grants.

The CBT seat allocation rule

The CBT Act assigns seats through a nomination structure that reflects a clear representational logic: every community that bore the costs of the Treaty should have a voice in how the revenues are spent. The current Act specifies a twelve-member board. Five seats are each nominated by a different regional district. One seat is nominated by the Ktunaxa Nation Council, the single First Nations body recognised in the Act. Six seats are appointed directly by the Province with no specific district nomination. The Lieutenant Governor in Council makes all final appointments and may decline nominees, meaning the Province retains formal appointment authority over the entire board. All twelve directors must reside within the Columbia Basin regardless of who nominates them.²

The Ktunaxa seat deserves attention. The CBT is structured as a regional organisation, but the inclusion of a Ktunaxa Nation seat gives it a different character than a purely regional institution. The Ktunaxa are not a regional district. They are a First Nation whose traditional territory covers a substantial portion of the Columbia Basin. Their inclusion means the CBT is not only a territorial compensation mechanism but also, implicitly, an inter-governmental arrangement between the Province and an Indigenous nation. The founding directors included two Ktunaxa representatives, reflecting the Ktunaxa/Kinbasket Tribal Council's active participation in the negotiations that preceded the 1995 Act. In the contemporary context, the Ktunaxa hold one seat, while four other Nations whose traditional territories overlap the Basin, namely Lheidli T'enneh, Secwépemc, Sinixt, and Syilx, hold none. This is an unresolved tension in the current governance structure but one that is outside the scope of this paper.

The board was substantially larger in its early years. Before the 2003 amendment, the board had approximately eighteen members, with a larger regional component and the Ktunaxa/Kinbasket Tribal Council holding two nominated seats rather than one. Under the original Act (Columbia Basin Trust Act, 1995), regional districts had the power to appoint their own directors directly, reflecting the communities' own negotiated settlement. In May 2001, the BC Liberal Party under Gordon Campbell won the provincial election and immediately subjected all provincial organisations to a core services review aimed at reducing expenditure and restructuring government. The CBT was included in this review (Columbia Basin Trust, 2020). In 2003, the Liberals passed Bill 79, the Columbia Basin Trust Amendment Act (Columbia Basin Trust Amendment Act, 2003), which reduced the board from eighteen to twelve members and replaced direct community appointment power with nomination power. Regional districts could now nominate candidates, but the Lieutenant Governor in Council would make final appointments and could decline any nomination. The restructuring also reduced the regional component from eleven to five seats and the Ktunaxa Nation from two seats to one. The one-seat-per-district principle, which is essential to our theory, remained unchanged.

What matters for testing our hypotheses is the structural consequence of the seat allocation rule. The same five regional districts plus the Ktunaxa Nation that each nominate one director also

²To our knowledge, no nominee has ever been declined.

compete for the board’s community grant budget. Any director who attempts to steer disproportionate grant funding homeward is simultaneously taking money from the nominating principals of five other regional directors plus the Ktunaxa nominee. Those directors face the same incentive to redirect funds homeward, not toward the first director’s community. A director from one community cannot form a majority by recruiting other regional directors. Each of them would prefer to steer money to their own community rather than support a rival’s claim. Provincial appointees have no incentive to join any regional extraction coalition. They are drawn from the Basin but answer to no specific district, and they constitute a permanent swing bloc that cannot be recruited into district-specific schemes.

In the next sections we outline our three empirical strategies. First, we use the synthetic control method to estimate the aggregate effect of the CBT endowment on regional economic development. Second, we use a random-effects model of grant allocation to test whether districts represented by politically experienced or long-tenured directors receive disproportionate funding. Third, we draw on seven waves of resident survey data (2010–2022) to examine whether Basin residents perceive the allocation as fair across the region.

4 Synthetic Control Design

We use the synthetic control method to construct a counterfactual trajectory for the Kootenay development region (Abadie and Gardeazabal, 2003; Abadie, 2021). The method builds a weighted combination of untreated units that best matches the treated unit on pre-treatment covariates and outcome trajectories. Post-treatment divergence between the treated unit and its synthetic counterpart is interpreted as the causal effect of the treatment. The method is well suited to settings with a single treated unit and an infrequent treatment event, both of which apply here.

Development regions are regional entities used for statistical analysis in Canada, grouping whole census divisions. British Columbia has eight development regions. The Kootenay development region almost entirely overlaps with the Columbia Basin area, encompassing the Regional Districts of Kootenay Boundary, Central Kootenay, and East Kootenay. Some Basin communities, including Revelstoke, Golden, and parts of the Columbia Shuswap, fall outside the Kootenay development region and are omitted from the outcome measure. This means we likely underestimate rather than overestimate the CBT’s developmental effect. We treat the \$250M endowment as a one-time windfall shock, ignoring the \$2 million in annual BC Hydro grants paid between 1990 and 1995, interest earnings, and returns on CBT investments. This further understates the total resource injection into the Basin over the study period.

Historical data for B.C. development regions are patchy prior to 1990. We address missingness using multiple imputation with Amelia, creating 1,000 imputed datasets and running the synthetic control on each. We present the full distribution of gaps rather than a single point estimate, capturing both imputation and estimation uncertainty. Our time period is $T = 1986\text{--}2019$ and the treatment event occurs at $T_0 = 1995$.

Data validation

Our main outcome of interest is local economic development. Statistics Canada does not release sub-provincial GDP data. As a proxy, we measure economic development using the number of business incorporation registrations per development region. British Columbia Statistics publishes annual incorporation counts by development region from 1990 onwards (BC Stats, 2024a). As a robustness check, we also gathered data on annual building permits, on the assumption that growth and development are positively related to business starts and construction (BC Stats, 2024b). Figure 1 shows the strong positive association between incorporations and annual building permit values across development regions ($r = 0.99$, $p < 0.001$).

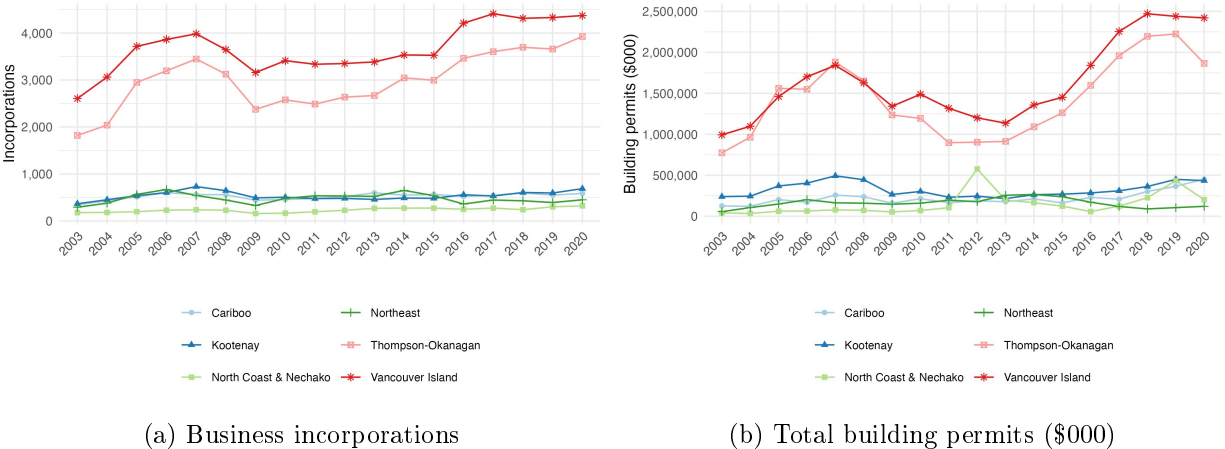


Figure 1: Business incorporations and total building permits by B.C. development region, 2003–2020. Lower Mainland suppressed for visual clarity but included in correlation analysis. Pearson $r = 0.99$ ($p < 0.001$).

Covariate balance and donor weights

The synthetic control uses six predictor variables measured over the pre-treatment 1986–1995 period. Three capture labour market conditions: unemployment rates in the tertiary, goods-producing, and service-producing sectors. Two capture employment levels: employment in the tertiary and goods-producing sectors. The sixth is the male population, which proxies for the working-age population in resource-sector economies where male employment dominates. These variables reflect the rural, resource-dependent character of the Kootenay economy and are the best available predictors of incorporation rates at the development region level over this period.

Table 1 shows that the match on these characteristics is close. The treated and synthetic units are similar on unemployment and employment across sectors, and the male population figures are nearly identical. The Lower Mainland, which is the most urban part of the province, contributes zero weight to the synthetic unit (Table 2). This is not imposed by the researchers. The algorithm assigns zero because the Lower Mainland does not resemble Kootenay on the pre-treatment trajectory or covariates. The donor pool is dominated by the Cariboo (0.30) and North Coast (0.47) regions,

which share Kootenay’s rural, resource-sector profile. The donor weights are consistent across all 1,000 imputed datasets.

Variable	Treated	Synthetic	Sample mean
Unemployment Tertiary	43.59	44.34	51.57
Unemployment Goods	61.32	107.42	98.54
Unemployment Services	50.93	43.67	50.67
Employment Tertiary	388.58	389.73	1,844.21
Employment Goods	289.60	158.43	653.31
Male population	70,611	70,823	266,419

Table 1: Covariate balance: treated and synthetic units.

Region	Weight
Vancouver Island	0.06
Lower Mainland	0.00
Thompson-Okanagan	0.05
Cariboo	0.30
North Coast	0.47
Nechako and North East	0.13

Table 2: Donor unit weights.

Results

Figure 2 shows the main results. The pre-treatment gap distributions sit tightly around zero, confirming good pre-treatment fit across the 1,000 imputed datasets. After 1995, the distributions shift upward. The estimated effect remains positive through to approximately 2013 before narrowing back toward zero. This inverted-U pattern is consistent across the full distribution of imputed datasets.

The timing is informative. A pure wealth effect from the CBT endowment would predict a flat positive shift that persists indefinitely. The data show a rise-and-fall pattern instead, with the positive gap peaking around 2001–2013 before narrowing back toward zero. This is consistent with the first scope condition for MTSA: the mechanism is most constraining when the community grant budget is genuinely fixed and competition among districts is acute. As total CBT benefits grew substantially after 2014, from approximately \$22 million to nearly \$100 million by 2019, the zero-sum character of community grant allocation weakened, and with it the structural stalemate that constrains home bias. The synthetic-control estimates therefore track the period of tightest budget constraint, not some confounding regional trend.

Over the full period, the CBT added approximately 560 incorporations in total. This is equivalent to roughly 1.3 years of average annual business registration activity in the Kootenay region.

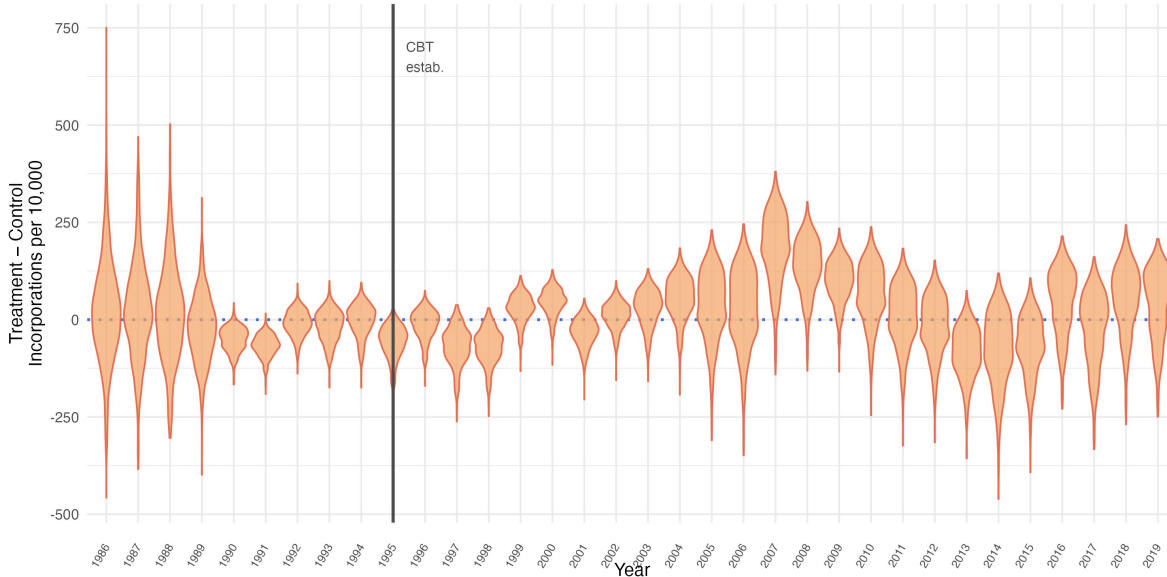


Figure 2: Distribution of synthetic control estimates across 1,000 imputed datasets. Each violin represents the distribution of $\hat{\tau}_{1t}$ for one year. Vertical line marks CBT establishment (1995). Pre-treatment distributions confirm parallel trends. The post-treatment inverted-U pattern is most pronounced 2001–2013.

At a community grant investment of \$45 million, this implies a cost of approximately \$80,000 CAD per additional business created. These are modest effects. The CBT is no panacea for regional economic development, but we do find positive effects from an institution that the resource curse literature would predict to generate nothing.

Hypothesis 1 predicted a null or negative effect consistent with the resource curse. The data reject this prediction across the full distribution of imputed datasets, though the magnitudes are small.

Robustness

The main result rests on a single treated unit and a small donor pool of six British Columbia development regions. We address this in four ways. First, the main figure incorporates uncertainty across 1,000 multiply-imputed datasets. The inverted-U pattern is consistent across the full distribution rather than being specific to any single imputed dataset. Second, we re-estimate the synthetic control dropping each donor region in turn (leave-one-out). The inverted-U pattern is consistent across all six leave-one-out specifications. The Lower Mainland contributes zero weight in the main analysis and its removal has no effect. Third, we backdate the treatment to 1990, when the provincial government and BC Hydro increased grant payments to Basin communities following the 1989 consultation described in Section 3. If the 1990 grant expansion rather than the 1995 CBT endowment were driving the positive gap, we would expect a positive effect to emerge around 1990. The 1990–1995 gap distributions sit close to zero. There is some indication of a peak in 2007

in this specification, but most years between 1990 and 2000 are not distinguishable from zero, inconsistent with the alternative interpretation that the pre-CBT grants drove later business growth. Fourth, we cross-check the incorporations measure against annual building permits, available across B.C. development regions from 2003 onwards. The two series are strongly correlated within the post-treatment window ($r = 0.99$ in Figure 1), supporting incorporations as a valid proxy for local economic activity in the periods where both can be observed.

The small donor pool limits the scope for alternative synthetic control specifications. Inference methods that require either a large number of donor units or equal donor weighting do not perform well in this setting. The four checks above are therefore the primary evidence that the main result is not an artefact of the estimation procedure, the imputation strategy, or the composition of the donor pool.

5 Grant Allocation and Board Composition

The synthetic-control estimate establishes that the CBT had a small, positive, cumulative effect on regional economic development. This section turns to the distributional question. Does the CBT board allocate grants equitably across the six nominating bodies, or do directors steer disproportionate funding toward their own communities?

Data

The grant data come from the CBT’s public project database, covering 1999–2020. We match each grant to the regional district where the recipient community is located, then sum to district-year totals. Our dependent variable is community-specific grants. We exclude Basin-wide grants, which are by construction distributed equally across all six districts and so cannot vary with director characteristics.

Director-level data come from a roster compiled from CBT annual reports (Columbia Basin Trust, various) and provincial Orders in Council. At any given time, roughly a third of sitting directors have prior or concurrent electoral experience, ranging from sitting members of the Legislative Assembly to municipal councillors and mayors. Across the full director-year panel, the share is higher, because politically experienced directors tend to serve longer tenures, so the panel is weighted toward their years of service. Among regional director-years specifically, 88 per cent are coded as having political experience. We return to the implications of this for identification below.

Figure 3 plots community grant funding by district from 1999 to 2020 in total and per capita. No district dominates grant receipts consistently. The lines cross repeatedly, with different districts receiving the largest allocations in different years. The Ktunaxa Nation, with a registered membership of approximately 1,381 across its four bands (BC Assembly of First Nations, 2024), has by far the smallest population of the six nominating bodies. On a per-capita basis, the Ktunaxa Nation’s total receipts (including Basin-wide programmes whose equal-share distribution across the six nominating bodies divides over a much smaller denominator) are several times higher than those

of any regional district. On community-specific grants alone, which are the variable that enters the regression analysis, the Ktunaxa rank in the middle of the distribution rather than at the top or bottom. The descriptive fingerprint that motivates the more formal analysis below is that no district sits systematically at the top of the community-specific grant distribution.

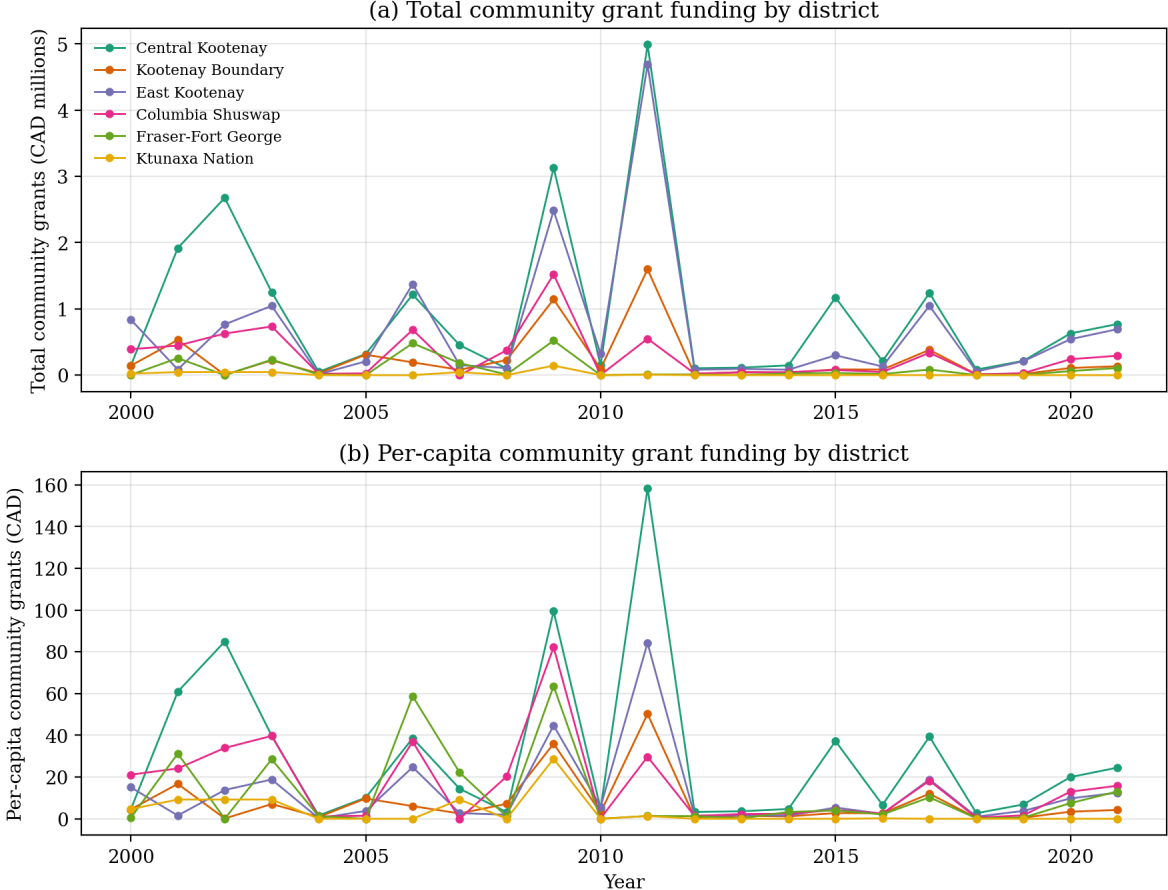


Figure 3: Community grant funding by district, 1999–2020. Top panel: total annual grant funding. Bottom panel: per-capita grant funding. No district dominates grant receipts consistently over time. Population denominators: RDCK 31,500; RDKB 31,800; RDEK 55,600; CSRD 18,500; FFGRD 8,200; Ktunaxa Nation 1,381.

Specification

We estimate a random-effects model linking annual grant allocations to two characteristics of the directors representing each district: their political experience and their tenure on the board.³ Polit-

³We use random rather than fixed effects on the district dimension because political experience is largely time-invariant within most districts: two of the six districts have constant political-experience coding across all years, and a third has only minor variation. A two-way fixed-effects estimator would absorb almost all the variation in PE through the district dummies, leaving only a thin slice of the data to identify β_1 . Random effects use both within- and between-district variation. Year fixed effects are retained throughout. Standard errors are clustered at the district level.

ical experience is coded as one for directors with confirmed prior, concurrent, or subsequent service as mayor, councillor, regional district representative, or provincial-level politician. Tenure length is the cumulative years served up to year t .

The model is

$$\text{Grants}_{i,t} = \beta_1 \text{PE}_{i,t} + \beta_2 \text{Tenure}_{i,t} + \beta_3 (\text{PE}_{i,t} \times \text{Tenure}_{i,t}) + \gamma_t + u_i + \varepsilon_{i,t},$$

where $u_i \sim N(0, \sigma_u^2)$ is the district random effect, γ_t are year fixed effects, and $\varepsilon_{i,t}$ is the idiosyncratic error. The random-effects assumption is that u_i is uncorrelated with the regressors. Given the descriptive evidence in Figure 3 that no district sits systematically at the top or bottom of the grant distribution, this assumption is defensible here.

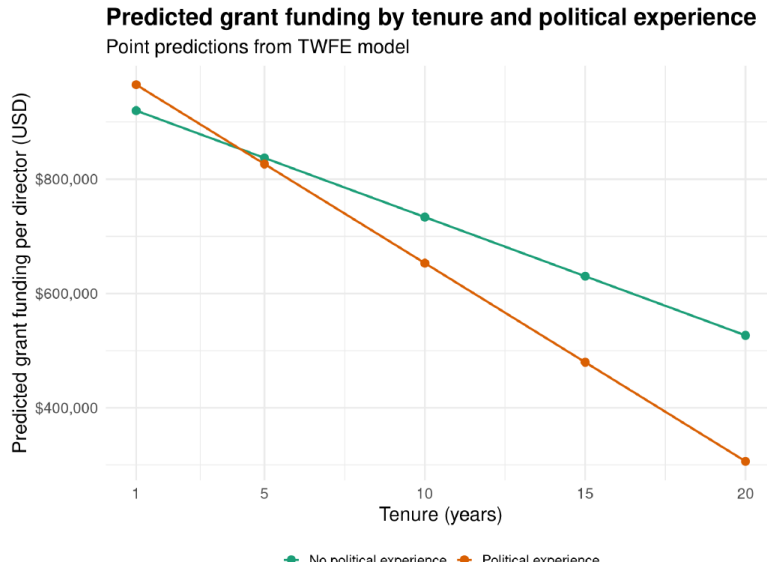
Table 3: Random-effects estimates of home-district grant allocation. Dependent variable: community-specific grants in CAD thousands. All specifications include year fixed effects and a district random effect estimated by Swamy-Arora. Cluster-robust standard errors at the district level in parentheses. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

	(1)	(2)	(3)
	PE only	+ Tenure	+ Interaction
Political experience	-19.7 (160.0)	-168.4 (195.1)	67.7 (136.5)
Tenure (years)		18.3 (13.2)	74.1* (43.1)
PE \times Tenure			-68.9 (44.4)
Year FE	Yes	Yes	Yes
District RE	Yes	Yes	Yes
Observations	132	132	132

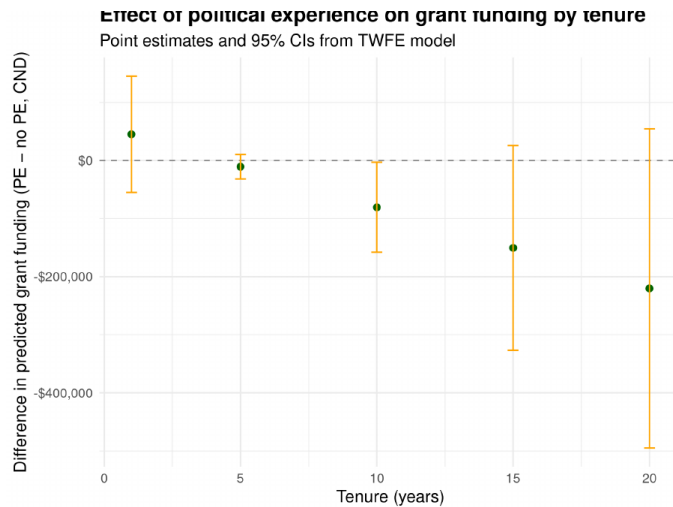
Results

Figure 4 reproduces the main random-effects estimates. Panel (a) shows predicted home-district grant funding by tenure for directors with and without prior electoral experience. The two lines diverge with tenure: predicted grants decline with tenure for both groups, but more steeply for politically experienced directors. Panel (b) shows the implied difference. At low tenure (1–5 years), the difference between politically experienced and inexperienced directors is small and positive. As tenure grows, the difference turns negative, with longer-tenured directors who have prior electoral experience associated with lower home-district grants than their less experienced counterparts.

Table 3 reports the random-effects estimates. The point estimates tell a consistent story. In the additive specification (PE and tenure separately), districts represented by politically experienced directors receive on average \$168,000 less per year in community-specific grants than districts without (cluster-robust SE \$195,000), and each additional year of tenure is associated with \$18,000 more in home-district grants (SE \$13,000). Neither coefficient is statistically distinguishable from



(a) Predicted grant funding by tenure and political experience



(b) Difference (PE minus no PE) by tenure

Figure 4: Predicted home-district grant funding by tenure and political experience. Panel (a): point predictions from the random-effects model. Panel (b): difference between politically experienced and inexperienced directors at various tenure values, with 95% confidence intervals. Cluster-robust standard errors at the district level. Year fixed effects.

zero. In the specification with the political-experience-by-tenure interaction, the interaction term is $-\$69,000$ per year of tenure (SE $\$44,000$, $p \approx 0.12$), implying that the gap between politically experienced and inexperienced directors widens by approximately seventy thousand dollars for each additional year a politically experienced director remains on the board. The signs run against home bias rather than with it.

Read at face value, this is consistent with the absence of home bias predicted by Hypothesis 2.

Politically experienced directors are the people most capable of building extraction coalitions, and they are associated with lower rather than higher home-district grants. The pattern strengthens with tenure, suggesting that institutional socialisation may compound the structural constraint over time.

What the data can and cannot establish

The result above should be read with caution. Three features of the data limit how strongly Hypothesis 2 can be tested.

First, the panel is small. It covers six districts (RDCK, RDKB, RDEK, CSRD, FFGRD, and the Ktunaxa Nation) over twenty-two years (1999–2020). Six clusters is a small number for cluster-robust inference, and the asymptotic justification for our standard errors should be read with that in mind.

Second, within-district variation in political experience is limited. Among regional director-years in the panel, 88% are coded as having political experience. Several districts have constant political-experience rates over the entire period (the Ktunaxa Nation seat is coded as zero throughout; East Kootenay’s Regional seat is coded as one throughout). Two districts (CSRD and FFGRD) account for most of the within-district variation. Political experience is also correlated with tenure ($r = 0.36$ among regional director-years; $r = 0.65$ in the full panel): politically experienced directors serve a mean of 10.0 years on the board, compared with 4.5 years for those without prior electoral experience. The two regressors share variance, which is why the interaction specification produces a larger negative point estimate even though neither main effect on its own is significant.

Third, the headline coefficients on political experience and the political-experience-by-tenure interaction are not statistically distinguishable from zero at conventional levels. The 95% confidence intervals on the predicted differences widen sharply with tenure (Figure 4b). Point estimates point in the direction predicted by Hypothesis 2, but the data do not rule out a null effect.

We also note that the structural mechanism described in Section 2 makes a sharper prediction than this panel can test. The mechanism is about how the *combined* board, comprising six district-rivals plus six provincial appointees with no district stake, reaches equilibrium. Our panel is at the district-year level. Provincial appointees are not attached to a district and so do not appear as a covariate in this design. The district-year regression can show whether districts represented by politically experienced nominees receive disproportionate funding (they do not). It cannot directly test whether the provincial bloc’s presence is what sustains the equilibrium, because the design holds the board composition fixed across all observations.

The honest reading is therefore the descriptive one. Across the six districts and twenty-two years, no district dominates grant receipts (Figure 3). The random-effects estimates point in a direction consistent with Hypothesis 2 but cannot reject the null of zero home bias from the data alone. We treat this as supporting evidence for the MTSA mechanism rather than as a direct test of it.

Within-district selection check

A natural concern is that politically experienced directors might be selected by communities expecting larger grant flows, a stepping-stone story in which connected districts place high-experience directors to capture more resources. We check this by comparing average political-experience rates and grant allocations by district across the pre- and post-2003 board structures. The within-district correlation between changes in political-experience rates and changes in grant allocations is 0.02. Districts where political-experience rates increased between periods did not receive more grants. This is inconsistent with the stepping-stone story and consistent with the equitable-allocation pattern documented in Figure 3.

6 Resident Perceptions: Demand-Side Evidence

The grant-allocation analysis offers supply-side evidence on how the CBT board distributes funding. We supplement this with demand-side evidence from seven biennial resident surveys conducted between 2010 and 2022, covering the full Basin population across four geographic quadrants. The 2010 and 2012 waves were conducted by Synovate/Ipsos. The 2014–2022 waves were conducted under the CBT’s standard resident survey programme. The questionnaire is consistent across waves. All data are CBT-commissioned, which likely inflates absolute satisfaction levels through social desirability bias. We therefore focus on cross-sectional and over-time *variation* rather than absolute levels.

The four quadrants are the CBT’s own community-liaison reporting units (Northeast, Northwest, Southeast, and Southwest), used in resident surveys throughout the period. The CBT staffs Community Liaisons by quadrant, with offices in Golden (Northeast), Nakusp (Northwest), and Cranbrook (Southeast). The quadrants do not correspond one-to-one to the six grant-allocation nominating bodies used in our supply-side analysis. The CBT does not publish a formal mapping of quadrants to regional districts, and the geographies are not nested. The two measure different things: the spending-side regression asks whether districts represented by certain kinds of directors receive more or less, and the survey asks whether residents in different parts of the Basin perceive the CBT differently. They are complementary rather than directly comparable, and we treat the demand-side evidence as a separate descriptive check rather than a within-unit confirmation of the regression result.

Survey participants

Across the seven waves, the total participant pool is 5,658 Basin residents. Table 4 summarises the design. Sample size per wave is approximately 800. Awareness of the CBT, the share of respondents who recognise the organisation when prompted, rose from 56% in 2010 to 81% in 2022. The familiar subsample (residents reporting at least some familiarity with the CBT’s role or history) is the base for the positive-difference questions used below. Familiar-subsample sizes range from 535 (2018) to 702 (2012).

Year	Total N	% aware of CBT	Familiar N
2010	811	56%	654
2012	815	59%	702
2014	808	—	660
2016	808	63%	644
2018	808	80%	535
2020	808	87%	590
2022	800	81%	573
Total	5,658	—	4,358

Table 4: Survey sample sizes per wave. “Familiar” subsample includes respondents who report at least some familiarity with the CBT’s role or history. The 2010–2012 waves were conducted by Synovate/Ipsos; the 2014–2022 waves were conducted under the CBT’s standard resident survey programme.

A 2022 microdata file containing individual-level responses, demographic covariates, and CBT-supplied survey weights is held by the project. Individual-level analysis using these microdata is reported in the working paper version available from the author. We do not reproduce the regression results here.

Are residents satisfied with the CBT?

Figure 5 presents two pieces of evidence. The left panel shows the percentage of familiar residents who agree the CBT is making a positive difference in their community and in the region as a whole, across all seven survey waves. The trend is consistent and high throughout. Community-level positive ratings range from 74% (2010–2012) to 87% (2018), stabilising at 82–84% in 2020–2022. The rise from 74% in 2010 to 87% in 2018 broadly tracks the inverted-U pattern in the synthetic-control estimates, suggesting that the CBT’s developmental impact and its perceived legitimacy moved together.

The right panel addresses the question directly relevant to the distributional mechanism. If home bias were occurring, with some districts systematically receiving less while others captured disproportionate shares, we would expect residents of disadvantaged districts to perceive the CBT less favourably. No such pattern is visible. Across all seven waves and all four quadrants, perceived positive difference in community ranges from 70% to 90%, with no quadrant systematically below the others. The cross-quadrant range within any single wave is at most ten percentage points (2016: Southeast 81%, Northeast 87%).

The cross-tabular pattern in Figure 5 carries the demand-side claim on its own. Sustained legitimacy at this level over twelve years is unusual for a government-affiliated body in a resource-extraction region. Three features of the data are worth commenting on.

The level is high. Across the seven waves, between 74% and 87% of familiar respondents agreed that the CBT makes a positive difference in their community. A standard concern with CBT-




Figure 5: CBT Resident Survey evidence, 2010–2022. *Left*: Percentage of familiar residents who agree the CBT is making a positive difference in their community (solid) and in the region as a whole (dashed). *Right*: Regional breakdown of community-level positive rating by quadrant. No quadrant is systematically below the others. Cross-quadrant range within any single wave is at most 10 percentage points.

commissioned surveys is social desirability bias: respondents may report more favourable views to interviewers asking on behalf of the institution being evaluated. We take this concern seriously and therefore interpret the level cautiously. The relevant comparison is not the absolute level but the relative pattern across waves and across quadrants, both of which should be unaffected by social-desirability inflation if the bias operates uniformly.

The trend tracks the developmental returns. The community-level positive rating rose from 74% in 2010 to 87% in 2018 and then fell back to 82% in 2022. The synthetic-control estimates show an inverted-U over a similar period, peaking around 2007–2013 before narrowing back toward zero. The two trajectories are consistent with the developmental impact and the perceived legitimacy of the CBT moving together. We are cautious about attributing causation, particularly because the surveys begin only in 2010, after the SC peak. But the descriptive correspondence is suggestive: residents perceived the CBT most favourably in the period when its measurable economic effects were largest.

Cross-quadrant variation is small. If home bias were occurring, with some communities systematically receiving less while others captured disproportionate shares, we would expect residents of disadvantaged areas to perceive the CBT less favourably. No such pattern is visible. Across all seven waves and all four quadrants, perceived positive difference in community ranges from 70% to 90%, with no quadrant systematically below the others. The cross-quadrant range within any single wave is at most ten percentage points (2016: Southeast 81%, Northeast 87%; 2010: Southeast 70%, Southwest 78%). The pattern complements the supply-side random-effects result rather than confirming it directly: residents in different parts of the Basin perceive the institution as

roughly equally fair, even where the four-quadrant survey geography only partially overlaps with the six-body grant-allocation geography.

7 Discussion

The MTSA mechanism predicts three things at the Columbia Basin Trust: a positive but bounded developmental return, the absence of systematic home bias in grant allocation, and resident perceptions of fairness across districts. The synthetic-control estimates confirm the first prediction. The CBT added approximately 560 incorporations to the Kootenay region between 1995 and 2019, with the effect concentrated in the period when the community grant budget was tightest (2001–2013). Resident surveys confirm the third prediction. No quadrant of the Basin perceives itself as systematically disadvantaged across seven waves of data spanning twelve years.

The second prediction, of no home bias in grant allocation, receives qualified support. The descriptive evidence is unambiguous: no district dominates grant receipts in any consistent way, and the lines cross repeatedly across the 1999–2020 period. The random-effects estimates point in the direction predicted by the theory: politically experienced and longer-tenured directors are associated with lower, not higher, home-district allocations. We are careful, however, not to overstate this. Within-district variation in director political experience is thin. The standard errors on the political-experience coefficient and its interaction with tenure are wide enough that the data do not rule out a null effect.

The structural mechanism we propose is about board-level equilibrium behaviour. Our district-year panel can show that districts represented by politically experienced or long-tenured nominees do not receive disproportionate funding, and they do not. The panel cannot directly identify the role of provincial appointees in sustaining the equilibrium, because their presence is constant across all observations. A more demanding test would require a research design that varies board composition or appointee type across institutions. The CBT alone cannot provide it.

Two secondary mechanisms may reinforce the structural constraint. The first is institutional socialisation. The CBT was created through genuine community mobilisation, and its founding documents stress equitable long-term stewardship. Long-serving directors are more likely to have absorbed these norms. The second is informal community accountability. The CBT operates in a small, close-knit region where visible favouritism would become known quickly. We cannot separate these mechanisms from the structural one with the data we have, but neither is required for the structural account to hold.

8 Conclusion

This paper examines whether internal board governance can substitute for electoral accountability in protecting resource windfalls from political capture. We argue that mixed-territorial seat allocation, when combined with zero-sum rivalry over a fixed budget, produces a structural stalemate that

prevents the formation of extraction coalitions. We test this using the Columbia Basin Trust, a benefit-sharing fund created in 1995 to compensate communities in British Columbia for the costs of reservoir construction under the Columbia River Treaty.

The CBT is both a likely and unlikely candidate for success. It lacks the two features the existing literature identifies as key: *ex ante* agreement on distribution rules and direct electoral accountability over the officials controlling the funds. Yet it operates in a high-capacity state with strong baseline institutions, and its board allocates grants through a structure that assigns one seat per competing community. We argue this raises the cost of home bias without requiring any of the external accountability infrastructure that standard remedies depend on.

Our results show modest but positive evidence of a cumulative effect on economic development: approximately 560 additional businesses created between 1995 and 2019, equivalent to roughly one additional year of regional growth stretched over a decade. The community grant budget of \$45 million translates to roughly \$80,000 CAD per additional business created, or about \$300 per Basin resident over the twenty-five-year period (\$45 million spread across approximately 147,000 residents). Standard political economy predicts null or negative returns from an unearned windfall managed by an unelected board. The CBT does not deliver such returns.

It is worth setting this against the comparison case the introduction motivates the paper around. Alaska’s Permanent Fund Dividend distributed \$1,702 per resident in 2024 and \$1,000 in 2025, with no measurable developmental effect (Goldsmith, 2010). The CBT’s per-person community-grant investment, roughly \$300 spread over twenty-five years, is approximately one-third of a single year’s Alaska PFD payout per Alaskan. For something on the order of a third of one year’s Alaska per-person return, the CBT delivered roughly one additional year of regional business growth and a structural arrangement that residents across all four quadrants perceive as fair. Whether this trade is worth making depends on what governments are trying to buy. If the goal is to put cash directly in residents’ hands, the PFD does that more efficiently. If the goal is to convert a windfall into productive local development without ceding discretion to electoral politics, the CBT design produces measurable returns that the PFD does not.

Random-effects models of grant allocation point to the absence of home bias. Politically experienced and long-tenured directors are associated with lower, not higher, home-district grant allocations. We are explicit that the political-experience coefficient is not statistically significant: within-district variation in this measure is limited, the panel is small, and the confidence intervals include zero. The strongest piece of evidence is the descriptive one. Across six districts and twenty-two years, no district consistently dominates grant receipts. Within-district correlations between changes in political-experience rates and changes in grant allocations are essentially zero, which is inconsistent with the stepping-stone story in which connected districts place high-experience directors to capture more resources.

The resident survey evidence reinforces this picture. Across seven biennial waves spanning 2010–2022, totalling 5,658 Basin residents, perceived positive difference in community ranges from 70% to 90% across all four quadrants. The supply-side absence of home bias and the demand-side absence

of perceived disadvantage line up.

Taken together, the three analyses produce a clear empirical pattern. The CBT delivers small positive developmental returns where the standard theory predicts none. It does not exhibit home bias. Residents across all four quadrants perceive the institution as fair. We are equally clear about what we cannot establish. We cannot show that MTSA itself, rather than institutional socialisation, founding norms, or informal community accountability in a small region, is the mechanism producing these patterns. Distinguishing between them would require a research design that varies board composition across institutions, which the CBT alone cannot provide. We treat the results as supporting evidence for the MTSA mechanism rather than as a direct test of it. What the data do not say is that MTSA fails. They say it is associated with patterns the resource curse literature predicts should not appear.

The central policy implication is specific. As governments create benefit-sharing bodies alongside hydropower, wind, and solar development, how seats are allocated across communities is a low-cost design choice with measurable consequences. A board where every regional seat represents a different community competing for the same budget raises the cost of coalition-building for extraction. This does not depend on the electoral infrastructure that alternative disciplining mechanisms require. Scotland's onshore wind community benefit framework, which sets a flat per-megawatt payment but leaves allocation to local boards, is one place where this design choice arises directly. The framework currently delegates allocation without specifying how seats on local boards should be allocated across competing communities. The CBT evidence suggests that this is precisely the design margin where institutional choice matters most.

Several questions remain open. The most pressing is how to identify MTSA's contribution separately from institutional socialisation, founding norms, and informal community accountability in a small region. We see three concrete research designs that would make progress on this.

The first is cross-institutional comparison. The MTSA prediction is that boards where every district-facing seat represents a different competing community produce equitable allocation, while boards that pool multiple districts under fewer seats or that lack a no-stake bloc do not. Building a panel of benefit-sharing institutions with varying seat structures, including hydropower funds in Norway, Quebec, and Brazil, and the Scottish and English onshore wind frameworks once they mature, would let one estimate the seat-allocation effect directly while holding constant the broader fiscal and political environment. The relevant variation is in the seat-to-district mapping and in the share of seats held by no-stake appointees. Both are observable and codeable from founding statutes.

The second is within-institution variation. Board structures change. The CBT itself reduced from eighteen to twelve seats in 2003, and the regional-to-no-stake ratio shifted in the process. Other benefit-sharing institutions have undergone similar amendments. A difference-in-differences design that exploits these changes, with grant allocation as the outcome and the seat-structure ratio as the treatment, would identify the MTSA contribution from within-institution variation rather than across-institution variation. The cost is that the timing of board reforms is rarely exogenous to the

broader political environment that may also affect allocation, but instrumental variable strategies based on legislative cycles or external core services reviews are available.

The third is direct measurement of board behaviour. Our paper observes the outcome of board decisions, namely grant allocations, but not the deliberative process producing them. Minutes would let researchers see whether the kinds of proposals that would produce home bias are tabled in the first place, modified before approval, or simply never raised. Votes are not a useful source on their own: contested votes are rare in CBT board practice, with most decisions reached by consensus, and a recorded objection is itself a rare event. The CBT minutes are not posted publicly, but the author has obtained the full set through a Freedom of Information request and is currently coding them in companion work. The interpretive payoff is to distinguish between coalition-formation attempts that fail, which the MTSA mechanism predicts, and boards that simply never attempt extraction in the first place, which would be consistent with norms-based explanations.

A fourth question is more substantive than methodological. The grant data distinguish community-specific from Basin-wide allocations but do not distinguish private from public goods within community grants. The MTSA mechanism predicts that public goods, projects benefiting multiple districts simultaneously, should be easier to fund than private goods targeted to a single district, because they reduce the zero-sum character of the allocation decision. Testing whether the composition of approved projects shifts with board structure or director tenure would provide a more direct test of the mechanism than aggregate allocation patterns alone.

Finally, this paper focuses on the distributional and developmental effects of windfall governance but does not examine effects on local politics. Two channels are plausible. CBT grants may affect who enters local politics: communities receiving more grant funding may attract more or different candidates for regional district seats, analogous to the political resource curse literature. Alternatively, grants may help incumbents stay in office by providing visible credit-claiming opportunities even in the absence of home bias in allocation. Both channels would connect the governance of resource windfalls to the broader political economy of rural communities and are testable with existing electoral data from British Columbia.

References

- Abadie, A. (2021). Using synthetic controls: Feasibility, data requirements, and methodological aspects. *Journal of Economic Literature*, 59(2), 391–425.
- Abadie, A., and Gardeazabal, J. (2003). The economic costs of conflict: A case study of the Basque Country. *American Economic Review*, 93(1), 113–132.
- Ahmed, F. Z. (2012). The perils of unearned foreign income: Aid, remittances, and government survival. *American Political Science Review*, 106(1), 146–165.
- Andersen, J. J., and Sørensen, R. J. (2022). The zero-rent society: Evidence from hydropower and petroleum windfalls in Norwegian local governments. *Journal of Public Economics*, 209, 104650.
- Baerg, N. (2020). *Crafting Consensus*. Oxford University Press.
- Baron, D. P., and Ferejohn, J. A. (1989). Bargaining in legislatures. *American Political Science Review*, 83(4), 1181–1206.
- BC Assembly of First Nations (2024). Kootenay region profile. <https://www.bcafn.ca/first-nations-bc/kootenay>.
- BC Stats (2024a). Annual business formations: Business incorporations by development region, 1990–2024. BC Data Catalogue, Government of British Columbia. <https://catalogue.data.gov.bc.ca/dataset/annual-business-formations>.
- BC Stats (2024b). Building permits monthly data. BC Data Catalogue, Government of British Columbia. <https://catalogue.data.gov.bc.ca/dataset/bc-stats-release-building-permits-bper->.
- Brollo, F., Nannicini, T., Perotti, R., and Tabellini, G. (2013). The political resource curse. *American Economic Review*, 103(5), 1759–1796.
- Columbia Basin Trust (2020). *Columbia Basin Trust: A Story of People, Power and a Region United*. Columbia Basin Trust.
- Columbia Basin Trust (various years). Annual Service Plan Reports. Columbia Basin Trust. <https://ourtrust.org/annualreport>.
- Columbia Basin Trust Act, S.B.C. 1995, c. 53 (now R.S.B.C. 1996, c. 53). Province of British Columbia.
- Columbia Basin Trust Amendment Act, 2003 (Bill 79). Province of British Columbia, 4th Session, 37th Parliament.
- Enikolopov, R. (2014). Politicians, bureaucrats and targeted redistribution. *Journal of Public Economics*, 120, 74–83.

- Fenno, R. F. (1978). *Home Style: House Members in Their Districts*. Little, Brown.
- Ferejohn, J. A. (1974). *Pork Barrel Politics*. Stanford University Press.
- Flavin, P. (2012). Income inequality and policy representation in the American states. *American Politics Research*, 40(1), 29–59.
- Fourinaies, A., and Mutlu-Eren, H. (2015). English bacon: Co-partisan bias in intergovernmental grant allocation in England. *Journal of Politics*, 77(3), 805–817.
- Goldsmith, S. (2010). The Alaska Permanent Fund Dividend: A case study in implementation of a basic income guarantee. Paper presented at the 13th Basic Income Earth Network Congress, University of São Paulo. Institute of Social and Economic Research, University of Alaska Anchorage.
- Havranek, T., Horvath, R., and Zeynalov, A. (2016). Natural resources and economic growth: A meta-analysis. *World Development*, 88, 134–151.
- Lillehammer, L. (2011). *Benefit Sharing and Hydropower Development: Glomma and Lågen Basin Case Study Report*. World Bank.
- Loeffler, H. (2023). Does a universal basic income affect voter turnout? Evidence from Alaska. *Political Science Research and Methods*, 11(3), 521–536.
- Mayhew, D. R. (1974). *Congress: The Electoral Connection*. Yale University Press.
- Nikolova, E., and Marinov, N. (2017). Do public fund windfalls increase corruption? Evidence from a natural disaster. *Comparative Political Studies*, 50(11), 1455–1488.
- Penfold, G. E. (2013). *A Review of the Range of Impacts and Benefits of the Columbia River Treaty on Basin Communities, the Region and the Province*. British Columbia Ministry of Energy, Mines and Natural Gas.
- Ross, M. L. (2001). Does oil hinder democracy? *World Politics*, 53(3), 325–361.
- Sachs, J. D., and Warner, A. M. (1995). Natural resource abundance and economic growth. NBER Working Paper No. 5398. National Bureau of Economic Research.
- Samson, C. (2016). Canada’s strategy of dispossession: Aboriginal land and rights cessions in comprehensive land claims. *Canadian Journal of Law and Society*, 31(1), 87–110.
- Schulz, C., and Skinner, J. (2022). Hydropower benefit-sharing and resettlement: A conceptual review. *Energy Research & Social Science*, 83, 102342.
- Shepsle, K. A., and Weingast, B. R. (1981). Political preferences for the pork barrel: A generalization. *American Journal of Political Science*, 25(1), 96–111.

Shugart, M. S., and Wattenberg, M. P. (eds.) (2001). *Mixed-Member Electoral Systems: The Best of Both Worlds?* Oxford University Press.

Yellowhead Institute (2021). Not so grand plans: The continued erasure of Indigenous rights in Newfoundland and Labrador's hydroelectric 'development'. Yellowhead Institute Brief.