

Incentive Alignment Bonds: A Mechanism Design Approach to Political Economy

Decentralized Institutes of Health

Mike P. Sinn

Abstract

Politicians systematically underinvest in public goods and overprovide concentrated benefits to organized interests. We propose Incentive Alignment Bonds (IABs): a mechanism that ties electoral support, post-office career opportunities, and legacy benefits to measurable outcomes on public good provision. By making politician utility functions depend on scores derived from their voting records on prioritized programs, IABs transform rent-seeking incentives into welfare-maximizing ones.

Table of contents

1	Abstract	3
2	Introduction	3
2.1	The Information-Incentive Disconnect	3
2.2	Definition	8
2.3	Why Reallocation, Not Addition	9
2.4	Political Change as an Asset Class	10
2.5	The Capital Asymmetry: Why IABs Can Outcompete Incumbent Lobbying	11
2.6	Quantifying Net Social Value of Funding Sources	12
2.7	Relation to Existing Literature	19
2.8	Contribution	19
2.9	Roadmap	20
3	Literature Review	20
3.1	Mechanism Design Theory	20
3.2	Public Choice Theory	21
3.3	Social Impact Bonds	21
3.4	Campaign Finance and Political Behavior	22
3.5	Credit Rating Agencies as Governance Mechanisms	22
4	Formal Model	23
4.1	Assumptions	23
4.2	The Politician's Utility Function	23
4.3	The Utility Function Transformation	24
4.4	Numerical Calibration and Parameter Ranges	25
4.5	The Policy Choice	27
4.6	The Pre-IAB Equilibrium	28
4.7	The IAB Mechanism	28

4.8	Score-Dependent Payoffs	28
4.9	Incentive Compatibility	29
4.10	Nash Equilibrium Analysis	30
4.11	Illustrative Example: A Global Health Treaty	30
4.12	Calibration: Parameter Ranges for Incentive Compatibility	32
5	The Three-Layer Architecture	32
5.1	Layer A: Scoring (Data Provision)	33
5.2	Layer B: Electoral Support	34
5.3	Layer C: Post-Office Benefits	34
5.4	Why This Is Not Bribery	34
5.4.1	Statutory Analysis: 18 U.S.C. § 201	34
5.4.2	Constitutional Protections	35
5.4.3	Established Precedent: Decades of Legal Scorecard Operations	35
5.4.4	The Key Legal Distinction	36
5.5	Legal Entity Separation	36
5.6	Funding Sources and Foundation Investment	36
6	Failure Modes and Limitations	37
6.1	Gaming and Metric Corruption	37
6.2	Plutocracy Objection and the “Lobbying With Extra Steps” Critique	37
6.3	Unintended Consequences	39
7	IABs in the Broader Governance Architecture	40
7.1	The Four-Layer Governance Stack	40
7.2	Why This Sequencing Matters	43
7.3	Comparison to Alternative Governance Mechanisms	43
7.4	What IABs Do Not Solve	44
8	Generalization: IABs for Global Coordination Problems	45
8.1	The General Template	46
8.2	Candidate Domains	46
8.3	Welfare Accounting of the IAB Mechanism	46
9	Conclusion	48
9.1	The Bootstrap Problem and Regulatory Resistance	49
	References	51
	Appendix A: Formal Proofs	58
9.2	Proof of Proposition 1 (Sufficient Condition for Incentive Compatibility)	58
9.3	Proof of Corollary 1 (Funding Threshold)	59
9.4	Proof of Proposition 2 (Multiple Equilibria Without IABs)	60
9.5	Proof of Proposition 3 (Equilibrium Selection With IABs)	60
	Appendix B: Detailed Application Specifications	61
9.6	B.1 Climate Change	61
9.7	B.2 Nuclear Disarmament	61
9.8	B.3 Pandemic Preparedness	62



Tip

Version 0.9 - Working Draft - Comments and critical feedback appreciated! Annotate via the sidebar or email mike@warondisease.org.

1 Abstract

This paper introduces Incentive Alignment Bonds (IABs), financial instruments that make supporting welfare-improving policies incentive-compatible for rational politicians. IABs are defined by three properties: (1) investor returns keyed to policy adoption, (2) politician benefits (electoral support, career prospects) based on voting record, and (3) funding from programs with lower net social value than the target use. Drawing on mechanism design theory and campaign finance research, we formalize a three-layer architecture (scoring, electoral, post-office) that transforms the politician's utility function: the same selfish optimization that currently favors concentrated interests can be redirected toward public goods.

We prove conditions for Nash equilibrium, provide explicit functional forms and numerical calibration showing feasibility with realistic parameters in low-opposition domains, and analyze legal constraints under U.S. anti-bribery law. The bootstrap problem (raising capital before regulatory opposition mobilizes) is solvable due to 272%+ expected ROI (from lobbying economics) and capital asymmetry: diffuse beneficiaries control \$454T vs. \$5T for concentrated interests. Under our global health treaty calibration, welfare accounting shows the mechanism's benefit-cost ratio exceeds 207:1 (Section 7.3). We demonstrate the mechanism through a detailed global health treaty application (Section 5), then analyze extensions to climate, nuclear risk, and pandemic preparedness (Section 8). Wishocracy (Layer 2) prevents allocation capture through democratic preference aggregation resistant to advertising-based gaming. We analyze failure modes, legal constraints, and implementation challenges in Sections 7-8.

The Problem in One Sentence

Politicians are currently rewarded for spending money on things that hurt society (weapons, fossil fuel subsidies) and punished for spending money on things that help society (cures, pandemic prevention), so government does the opposite of what would make our lives better.

The Solution: Incentive Alignment Bonds (IABs) flip the incentives by creating a huge pool of money that rewards politicians (with campaign support and cushy post-office jobs) for supporting the stuff that actually helps people.

2 Introduction

2.1 The Information-Incentive Disconnect

A central puzzle in political economy is why policies with large positive net social welfare often fail to be adopted. The economic case for many global public goods, including climate mitigation, pandemic preparedness, and clinical trials, is overwhelming, yet governments don't do them. **The**

Aligning Politics with the Public Good: An Introduction to Incentive Alignment Bonds

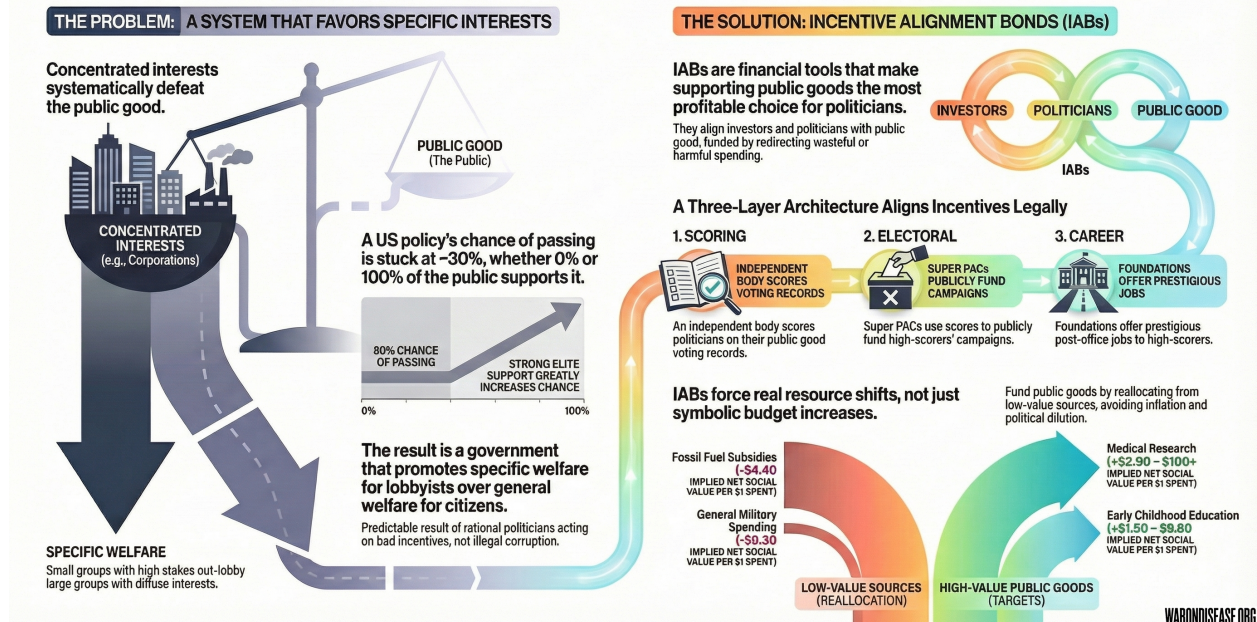


Figure 1: The Incentive Alignment Bonds Mechanism

conventional diagnosis is wrong: the problem is not lack of information or resources, but wrong incentives.

Rankings of government programs by net social value already exist and are systematically ignored. The Copenhagen Consensus has published rigorous benefit-cost ratio (BCR) analyses since 2004 (“Halftime for SDGs: Child Immunization” 2023). Their findings are clear: childhood vaccinations (101:1 BCR), e-government procurement (125:1), and maternal health interventions (87:1) dramatically outperform military spending beyond deterrence requirements (~0.7:1) and fossil fuel subsidies (negative net social value). GiveWell, Open Philanthropy, the IMF, and numerous academic institutions produce similar analyses.

Yet government spending patterns have not shifted. The U.S. spends \$16 on military operations for every \$1 on diplomacy and humanitarian aid. Programs with benefit-cost ratios exceeding 100:1 receive single-digit billions while programs with negative net social value receive hundreds of billions.

Put another way: If a private equity firm allocated capital like the U.S. government, it invests \$100M in ventures that destroy \$500M in value while rejecting ventures that return \$10B. The firm would be bankrupt within a year. The government simply prints more money and calls it “fiscal policy.”

The marginal value of producing another ranking is zero. We don’t need more information. Politicians already know which programs produce net social value. We need a mechanism to make existing information consequential in the political utility function (the mathematical representation of what politicians actually care about: reelection, money, status).

This paper argues that politicians maximize reelection probability, post-office career prospects, and status, not aggregate social welfare. **The problem is not that decision-makers don’t know which programs produce value. They do know. They don’t care, because caring doesn’t**

appear in their utility function.

Mancur Olson’s seminal work *The Logic of Collective Action* (1965) established that groups with concentrated interests (small groups with huge per-person stakes, like defense contractors) systematically outcompete groups with diffuse interests (millions of people with small individual stakes, like citizens who’d benefit from cures) in political competition. Why? A small group facing large per-member stakes will invest more in lobbying than a large group facing small per-member stakes. This holds even when total welfare losses exceed total welfare gains. The beneficiaries of inefficient policies (defense contractors, pharmaceutical incumbents, fossil fuel producers) are concentrated. The beneficiaries of efficient policies (citizens who would benefit from cures, climate stability, reduced existential risk) are diffuse.

The result is a systematic distortion of democratic governance. The U.S. Constitution’s Preamble charges government to “promote the general Welfare,” the welfare of all citizens, not the welfare of any particular faction. Yet Olsonian dynamics ensure that government routinely promotes *specific* welfare (of concentrated interests who can afford to lobby) at the expense of *general* welfare (of the diffuse public). This is not corruption in the legal sense; it is the predictable equilibrium (stable outcome where nobody wants to change their behavior) of rational actors operating within existing institutional rules. The problem is not bad actors but bad incentives.

Empirical evidence confirms this pattern. Gilens and Page (2014) analyzed 1,779 policy decisions over two decades. Finding: “economic elites and organized groups representing business interests have substantial independent impacts on U.S. government policy, while mass-based interest groups and average citizens have little or no independent influence.” The correlation between average citizen preferences and policy outcomes was effectively zero. Whether 0% or 100% of citizens supported a policy, its probability of adoption remained unchanged at approximately 30%. By contrast, policies favored by economic elites were adopted at significantly higher rates. Figure 2 illustrates this disconnect.

Key Results

This paper demonstrates that the adoption problem (getting welfare-improving policies passed despite concentrated opposition) is solvable through mechanism design (designing rules so people’s selfish choices create good outcomes):

- **272% expected first-year ROI** on bootstrap capital from lobbying economics (Section 2.2)
- **\$454T vs. \$5T capital asymmetry** ensures diffuse beneficiaries (citizens who’d benefit from cures) can outspend concentrated opposition (defense contractors) by 90:1 (Section 2.3)
- **207:1 benefit-cost ratio** of the IAB mechanism itself, even accounting for overhead costs and failure risk (Section 7.3)
- **Mechanism achieves incentive compatibility** (when doing the right thing is also the profitable thing) with realistic parameters in low-opposition domains; requires higher funding in entrenched-opposition domains (Section 3.3)
- **Wishocracy (Layer 2) prevents allocation capture** through random pairwise sampling that makes advertising economically infeasible (Section 6.2.3)
- **Bootstrap problem is solvable:** ROI potential attracts risk-tolerant capital; capital asymmetry ensures scalability once proof-of-concept succeeds (Section 8.1)




Figure 2: The relationship between public support and policy adoption. Whether a policy has 0% or 100% public support, its probability of adoption remains flat at around 30%. By contrast, policies favored by economic elites show a strong positive correlation with adoption. Data from Gilens and Page (2014) analysis of 1,779 policy decisions.

Reader's Guide

For non-technical readers: Read Sections 1-2 (problem and solution overview), 6 (governance architecture), and 7-8 (welfare analysis and implementation). The formal model (Section 3) can be skipped; the intuition is provided in Section 2 and visualized in figures.

For economists: Focus on Section 3 (explicit functional forms and numerical calibration), Section 5 (legal architecture), Section 7 (welfare accounting), and Section 8 (bootstrap solution). The calibration in Section 3.3 shows both failure and success cases with realistic parameters.

For investors and philanthropists: Read Section 2.2 (ROI calculation), Section 2.3 (capital asymmetry), Section 7.3 (mechanism BCR), and Section 8.1 (bootstrap problem solution). Expected returns derive from proven lobbying economics, not speculative outcomes.

Key Terms for Non-Economists (Click to Expand)

The Big Concepts:

- **Mechanism design:** Designing rules so people's selfish choices create good outcomes (like auctions that make bidders reveal true value)
- **Incentive compatibility:** When doing the right thing is also the profitable thing
- **Nash equilibrium:** A stable situation where nobody can improve by changing their strategy alone
- **Utility function:** Mathematical representation of what someone cares about (reelection, money, legacy)
- **Dominant strategy:** Best choice regardless of what others do

Political Economy:

- **Concentrated interests:** Small groups with huge per-person stakes (defense contractors losing billions)
- **Diffuse beneficiaries:** Millions of people with small individual stakes (each citizen gains 5 healthy years)
- **Olsonian dynamics:** Named after economist Mancur Olson. It is the pattern where small organized groups beat large disorganized groups in political competition
- **Collective action problem:** Hard for large groups to coordinate even when everyone would benefit
- **Public goods:** Things everyone benefits from (clean air, medical research, pandemic preparedness)
- **Regulatory capture:** When regulated industries control their regulators
- **Rent-seeking:** Lobbying for special favors rather than creating value

Finance & ROI:

- **ROI (Return on Investment):** Profit divided by cost (272% ROI means you get \$3.72 back for every \$1 invested)
- **Capital asymmetry:** Unequal distribution of wealth (\$454T households vs. \$5T defense contractors)
- **Securitization:** Turning something into tradable financial instruments (like turning mortgages into mortgage-backed securities)
- **Expected value:** Average outcome accounting for probabilities (20% chance of \$100 = \$20 expected value)
- **Market capitalization:** Total value of a company's stock

Game Theory & Economics:

- **Strategic agents:** People making calculated choices based on self-interest
- **Equilibrium outcome:** Stable result when everyone is doing their best given what others are doing
- **Ex ante / Ex post:** Latin for “beforehand” / “afterwards”
- **Quid pro quo:** Latin for “this for that.” It refers to a direct exchange

Domain-Specific:

- **Social Impact Bonds (SIBs):** Investors fund social programs, get repaid only if outcomes achieved
- **BCR (Benefit-Cost Ratio):** For every \$1 spent, how many \$ of social benefit (100:1 BCR = \$100 benefit per \$1 spent)
- **NSV (Net Social Value):** Total benefit minus total cost to society
- **DALYs:** Disability-Adjusted Life Years. This is a measure of disease burden combining mortality and quality of life
- **VSL (Value of Statistical Life):** How much society values preventing one death (EPA uses ~\$10M)

For additional terms: definitions in parentheses on first use throughout the paper.

This paper introduces **Incentive Alignment Bonds (IABs)**, a mechanism design approach to reversing this dynamic. Rather than attempting to change politicians’ preferences or relying on benevolent decision-makers, IABs restructure the incentive environment so that rational self-interest points toward welfare-improving policies.

2.2 Definition

An Incentive Alignment Bond is a financial instrument with three primitive properties:

1. **Investor alignment with public good production:** Investors receive returns proportional to verified public-good funding flows achieved (e.g., treaty ratification, appropriations enacted). Returns are keyed to observable funding events, not downstream outcomes, avoiding complex attribution problems. If the policy succeeds, investors profit; if it fails, they do not. This creates a class of actors with concentrated financial interest in policy success.
2. **Politician alignment with public good production:** Politicians receive electoral support and career benefits based on their voting record for the target policy class. Supporting the public good becomes the utility-maximizing choice, increasing reelection probability, post-office income, and status.
3. **Funding from lower-value sources:** The public good is funded by redirecting resources from government activities with lower social value than the target use. In the strongest case, these are programs that produce measurable net harm. More generally, they are programs that persist due to Olsonian concentrated-benefit/diffuse-cost dynamics rather than comparative merit. This constraint ensures IABs reallocate from less valuable to more valuable uses; they are welfare-improving in expectation under standard utilitarian social welfare assumptions.

The third property is crucial. Many government expenditures persist not because they produce social value, but because their beneficiaries are concentrated while their costs are diffuse. Military spending beyond deterrence requirements. Agricultural subsidies that distort markets. Fossil fuel subsidies that accelerate climate change. These programs survive political competition not on merit but on lobbying intensity. Concentrated interests (defense contractors, incumbent industries) invest

heavily in maintaining these programs, while diffuse beneficiaries face collective action problems. IABs redirect these resources to measurable public goods, making the reallocation welfare-improving even before accounting for the public good produced.

The key innovation is that all three alignments reinforce each other. A fraction of redirected funding flows perpetually to (a) investor returns and (b) political incentive mechanisms, while (c) the funding source ensures no beneficial programs are displaced. This makes the instrument self-sustaining: policy success generates the resources that maintain both investor and politician alignment, creating pressure for continuation and expansion.

Unlike Social Impact Bonds (which align service providers with program outcomes), IABs align the actors who control large-scale budget decisions, the politicians, with public-good production. Unlike lobbying (which aligns politicians with narrow interests), IABs align politicians with measurable, welfare-improving outcomes. Unlike campaign contributions (which are ad hoc and candidate-specific), IABs create universal, ex-ante rules tied to objective metrics. And unlike new taxation or deficit spending, IABs fund public goods by displacing harmful or wasteful expenditure.

A natural objection: “How is this not just PACs + voter scorecards + SIBs?” The answer is that existing mechanisms fail to solve the Olsonian problem because they operate in isolation. PACs exist but remain dominated by concentrated interests due to the 90:1 capital asymmetry. Scorecards exist (League of Conservation Voters, NRA ratings) but lack teeth without accompanying financial incentives. SIBs exist but require government to agree upfront, which faces the same collective action problem IABs solve. IABs integrate all three components (investor capital, electoral accountability, post-office incentives) into a single instrument where each component’s effectiveness depends on the others, and where the funding mechanism creates permanent financial incentives for diffuse beneficiaries to overcome collective action barriers. The whole is greater than the sum of its parts because the integrated architecture makes coordinated capital deployment economically rational for millions of individual investors.

2.3 Why Reallocation, Not Addition

A natural question: why insist on *redirecting* resources from harmful programs rather than simply *adding* new spending on public goods? The answer: real resources (people, factories, scientists) are finite.

When advocates successfully lobby for increased spending on a beneficial program, the political response is rarely a corresponding reduction in harmful programs. Instead, governments usually:

1. **Expand the overall budget** through deficit spending or monetary expansion
2. **Increase spending on politically powerful programs** in parallel, to maintain coalition support
3. **Dilute the real value** of new spending through inflation and competition for fixed resources

The result is that the beneficial program’s *nominal* budget increases while its *real* share of resources remains unchanged or even decreases. The concentrated beneficiaries of harmful programs lose nothing; the diffuse beneficiaries of the new program gain less than the headline figures suggest.

Consider the empirical pattern: between 2000 and 2024, U.S. federal spending on both defense and healthcare increased much in nominal terms. Neither constituency “lost” to the other. Instead, total federal spending grew from \$1.8 trillion to \$6.1 trillion, while the national debt expanded from \$5.6 trillion to \$34 trillion. The real constraint, the finite pool of engineers, scientists, manufacturers, and

productive capacity, was diluted across an ever-expanding set of nominal claims. Defense contractors retained their share of real resources; healthcare advocates won nominal victories that competed with inflation and expanding claims elsewhere in the budget.

This is not a failure of advocacy. It is the equilibrium outcome (stable result when everyone optimizes given current rules) of Olsonian dynamics operating under soft budget constraints (when governments can spend without immediate consequences; they just print money or charge it to the credit card of future generations). Concentrated interests protect their programs absolutely. Diffuse interests win nominal victories. The budget grows in nominal terms while each program's claim on real resources remains contested.

The mechanism is straightforward: when Congress appropriates \$10 billion for pragmatic clinical trials without specifying a funding source, the Treasury either borrows or the Federal Reserve accommodates. Total nominal spending increases.

But the supply of trained researchers, laboratory equipment, and institutional capacity does not increase proportionally. The new dollars compete with existing dollars, including those flowing to defense, fossil fuel subsidies, and other Olsonian programs, for the same finite resources.

Inflation, both general and sector-specific, erodes the real purchasing power of the nominal increase.

IABs break this dynamic by *specifying the funding source* as part of the mechanism. The policy does not say “fund pragmatic clinical trials”; it says “redirect 1% of military spending to pragmatic clinical trials.” This forces a real reallocation:

- Military budgets face *actual* cuts in nominal and real terms
- Concentrated losers (defense contractors) face *real* losses, which they will resist
- But the IAB political incentive layer ensures politicians who support the reallocation benefit more than those who resist
- The finite pool of resources shifts from net-negative to net-positive uses

The third primitive property, funding from harmful sources, is therefore not merely a normative preference or a political strategy. Under the soft budget constraint dynamics described above, it is enough to achieve real resource reallocation (likely needed in practice). Without specifying the funding source, advocates win symbolic victories while real resources continue flowing to Olsonian programs. With specified reallocation, the IAB creates genuine redistribution from net-negative to net-positive uses, constrained only by political will rather than by the illusion of unlimited budgetary capacity.

2.4 Political Change as an Asset Class

To understand why investors would fund the initial campaign to pass an IAB treaty, we must recognize political advocacy not as charity, but as a high-yield asset class. The Return on Investment (ROI) for corporate lobbying is historically orders of magnitude higher than traditional financial markets. The following illustrative estimates, drawn from existing studies, suggest order-of-magnitude returns:

- **Defense:** The top five defense contractors spent \$1.1 billion on lobbying over two decades while receiving \$2.02 trillion in Pentagon contracts, an implied ROI of approximately **181,000%** (Smithberger and Hartung 2021).
- **Fossil Fuels:** The oil and gas industry spent \$151 million lobbying in 2024 to protect \$17 billion in subsidies, an implied annual return of approximately **11,000%** (“Federal Lobbying Set New Record in 2024” 2025).

- **Medical Research:** Academic analysis finds that for specific disease groups, each \$1,000 spent on lobbying is associated with \$25,000 in additional NIH funding, an implied return of approximately **2,500%** (Reardon 2014).

These figures are illustrative rather than precise point estimates; causal attribution is complex and effects vary by context. Nevertheless, they suggest that political influence generates returns far exceeding traditional asset classes (compare to the S&P 500's historical average of ~10%). Political change is currently accessible only to concentrated industries protecting the status quo. IABs securitize this opportunity, allowing investors to capture a fraction of the value generated by shifting government priorities toward public goods.

$$E[R] = \frac{P(\text{success}) \times (\text{Redirected Flow} \times \text{Investor Share})}{\text{Campaign Cost}}$$

Illustrative calculation (base case): If a campaign costs \$100 million and has a 20% chance of passing a treaty that redirects \$24 billion/year (with 10% to investors), the expected annual value is:

$$E[V] = 0.20 \times \$2.4 \text{ billion} = \$480 \text{ million/year}$$

The expected first-year ROI is 272%, with the revenue stream continuing in perpetuity.

Sensitivity to parameter uncertainty:

Parameter	Conservative	Base Case	Optimistic	Expected First-Year ROI
Success probability	10%	20%	40%	36% / 444% / 2,076%
Campaign cost	\$200M	\$100M	\$50M	36% / 444% / 2,076%
Investor share	10%	10%	15%	36% / 444% / 2,076%
Treaty size	\$27.2B/year	\$24B/year	\$48B/year	36% / 444% / 2,076%

Even under conservative assumptions (10% success, \$200M cost, 10% investor share, \$27.2B treaty), expected first-year ROI is 36% positive. This is significantly higher than traditional asset classes. The calculation demonstrates that the mechanism is economically rational for risk-tolerant capital across a wide parameter range, not dependent on optimistic assumptions.

2.5 The Capital Asymmetry: Why IABs Can Outcompete Incumbent Lobbying

Here's what almost everyone misses: **the collective capital available to diffuse beneficiaries vastly exceeds the resources of concentrated interests.** The problem is not lack of resources but lack of coordination.

Consider the global capital distribution:

- **Concentrated interests** (defense contractors, fossil fuel companies, pharmaceutical incumbents): Combined market capitalization ~\$5T¹
- **Diffuse beneficiaries** (everyone who would benefit from cures, climate stability, pandemic prevention): Global household wealth ~\$454T (“Global Wealth Report 2023” 2023)

The 90:1 capital advantage cannot currently be deployed because diffuse beneficiaries face a collective action problem (hard for large groups to coordinate even when everyone would benefit): each individual’s stake is small, and coordination costs are prohibitive. A retiree who would gain 5 additional healthy years from accelerated medical research has enormous stake (~\$500K in value of statistical life-years) but cannot coordinate with millions of other retirees to match pharmaceutical lobbying budgets.

IABs solve the coordination problem by securitizing political change (turning policy outcomes into tradable financial instruments, like turning mortgages into mortgage-backed securities). If IABs can be structured as tradable securities that any individual can purchase, they transform diffuse beneficiaries into the largest special interest group in history. The mechanism:

1. **Investment accessibility:** Structure IABs as securities available to retail investors (similar to green bonds or social impact bonds)
2. **Returns exceed alternatives:** Expected returns of 100-1000%+ (from lobbying ROI) dramatically exceed stock market returns of 10%
3. **Massive capital mobilization:** Even 0.1% of global household wealth (\$454 billion) exceeds total annual global lobbying spending by 100x
4. **Self-interest alignment:** Investors profit directly from policy success, creating concentrated financial incentives on the welfare-improving side

This reverses the Olsonian asymmetry. Defense contractors spend \$100M+ annually lobbying because each firm captures concentrated benefits.

But the IAB mechanism allows millions of diffuse beneficiaries to collectively deploy billions while each capturing a proportional return.

The concentrated opposition (\$100M-\$1B annually) becomes outmatched by newly-coordinated diffuse support (\$10B-\$100B+ available capital).

The key insight: We’re not inventing lobbying. Defense contractors proved it works decades ago. We’re just making it available to everyone whose lives depend on cures instead of bombs. Turns out there are more of us, and we’re much richer. The political change ROI (100-10,000x) that currently accrues only to concentrated industries becomes accessible to everyone who benefits from public goods.

2.6 Quantifying Net Social Value of Funding Sources

The third primitive property requires identifying government expenditures with lower marginal social value than the proposed alternative. This is an empirical question that can be addressed using standard welfare economics:

¹Conservative estimate based on: Top 100 defense contractors ~\$2T market cap (Lockheed Martin \$120B, Raytheon \$100B, etc.); top fossil fuel companies ~\$2.5T (ExxonMobil \$400B, Saudi Aramco \$2T, Shell \$200B, etc.); top 20 pharmaceutical companies ~\$4T (but most align with IABs for medical research funding). Using \$5T assumes pharmaceutical opposition, which may not materialize. Actual concentrated opposition ranges \$2T-\$5T depending on domain.

$$NSV_i = B_i - C_i \quad (1)$$

where NSV_i is net social value of program i , B_i is measurable social benefits (including economic multiplier effects), and C_i is total costs (including both direct expenditure and opportunity costs of foregone alternatives).

A program is an appropriate IAB funding source if:

$$NSV_{source} < NSV_{target} \quad (2)$$

That is, if the marginal social value of the source program is lower than that of the target program. In the strongest case, $NSV_{source} \leq 0$; the program produces net harm. But reallocation is justified even when both programs have positive value, as long as the target exceeds the source.

This comparative criterion implies a natural ranking of government expenditures by marginal social value. IABs should draw from the *lowest-ranked* programs first, those that persist mainly due to Olsonian dynamics rather than merit. This avoids the objection that “defense isn’t ALL bad”; the claim is not that military spending is worthless, but that its marginal value is lower than medical research, and that it persists at current levels due to lobbying intensity rather than comparative social value.

Military spending provides a concrete example. The economic multiplier for military spending is around 0.6, compared to 1.5–3.0 for infrastructure, education, and medical research. This means each dollar of military spending generates \$0.60 in economic activity, while each dollar of medical research investment generates \$2.00–3.00. The opportunity cost alone, before considering direct harms, makes military spending beyond minimum deterrence requirements a net social loss.

More comprehensive analysis includes:

Table 2: Estimated Annual Societal Cost of Global Military Spending

Cost Category	Annual Value	Source
Direct military expenditure	\$2.4T	SIPRI 2024
Lost GDP from multiplier differential	\$2.4T	$(M_{alt} - M_{mil}) \times S$
Infrastructure destruction (active conflicts)	\$1.9T	World Bank
Human casualties (VSL method)	\$2.4T	EPA VSL \times conflict deaths
Trade disruption	\$0.6T	World Bank trade flow analysis
Veteran healthcare (ongoing)	\$0.2T	VA budget analysis
Environmental damage	\$0.1T	UNEP conflict assessments
Total societal cost	\$10T+	Sum of above

A large literature suggests the net social value of military spending beyond minimum deterrence is strongly negative: the \$2.4 trillion in direct expenditure produces \$10+ trillion in societal costs

(see Table 2 for illustrative order-of-magnitude estimates). If these estimates are roughly correct, redirecting even 1% to pragmatic clinical trials, which has positive externalities and high multipliers, would be welfare-improving on net.

Similar analysis applies to other Olsonian programs:

- **Fossil fuel subsidies** (\$7T globally, IMF 2023): Accelerate climate change, distort energy markets, produce negative externalities exceeding subsidy value
- **Agricultural subsidies** in developed nations: Distort global food markets, harm developing-country farmers, produce environmental damage from monoculture incentives
- **Regulatory capture programs**: Expenditures that exist to protect incumbent firms from competition rather than serve public interest

Table 3 presents a comprehensive ranking combining Copenhagen Consensus data with other authoritative sources. The pattern is stark: programs with BCRs exceeding 100:1 receive single-digit billions in annual spending, while programs with negative NSV receive hundreds of billions. **Spending correlates with lobbying intensity, not social value.** This systematic misallocation is the empirical foundation for the IAB mechanism.

Table 3: Ranking of Government Expenditures by Benefit-Cost Ratio and Net Social Value

Program / Category	BCR	Annual Spending	Est. NSV (per \$1)	Source
HIGH-NSV PROGRAMS (IAB TARGETS)				
Childhood Vaccinations (Global)	101:1	~\$8B	+\$100	Copenhagen Consensus 2023 (“Halftime for SDGs: Child Immunization” 2023)
E-Government Procurement	125:1	~\$2B	+\$124	Copenhagen Consensus 2023 (“Halftime for SDGs: Child Immunization” 2023)

Program / Category	BCR	Annual Spending	Est. NSV (per \$1)	Source
Maternal/Neonatal Care	87:1	~\$12B	+\$86	Copen- hagen Consensus 2023 ("Halftime for SDGs: Child Immuniza- tion" 2023)
Nutrition Interventions	18:1	~\$5B	+\$17	Copen- hagen Consensus 2023 ("Halftime for SDGs: Child Immuniza- tion" 2023)
Medical Research (NIH)	2.56:1 — 4.75:1	\$47B	+\$1.56 – \$3.75	United for Medical Research ("Annual NIH Economic Impact Report: 2025 Update" 2025)
Early Childhood Education	2.5:1 – 10.8:1	~\$30B	+\$1.50 – \$9.80	(Heckman et al. 2010), (Reynolds et al. 2011)
MODERATE-NSV PROGRAMS				

Program / Category	BCR	Annual Spending	Est. NSV (per \$1)	Source
Infrastructure (High-Quality)	1.5:1 – 2.5:1	Varies	+\$0.50 – \$1.50	(“Estimated Impact of the American Recovery and Reinvestment Act on Employment and Economic Output in 2014” 2015); highly project-dependent
LOW/NEGATIVE-NSV PROGRAMS (IAB SOURCES)				
Military (Beyond Deterrence)	~0.7:1	\$450B+ (US)	-\$0.30	(Barro and Redlick 2011); fiscal only, excludes conflict costs
Fossil Fuel Subsidies (Explicit)	<0	\$20B (US)	-\$1 to -\$5	(“IMF Fossil Fuel Subsidies Data: 2023 Update” 2023)
Fossil Fuel (Implicit/Externalities)	-5:1	\$568B (US)	-\$5	(Batini et al. 2021), PNAS 2021 externality analysis

The empirical criterion is straightforward: if a program’s beneficiaries are concentrated, its costs

are diffuse, and rigorous cost-benefit analysis shows $NSV_{source} < NSV_{target}$, reallocation is welfare-improving. In the clearest cases, where $NSV_{source} \leq 0$, the source program is actively harmful and reallocation is unambiguously beneficial. But even programs with modestly positive NSV are appropriate sources if the target has much higher value. The IAB mechanism does not require moral judgment about the source program, only comparative measurement.

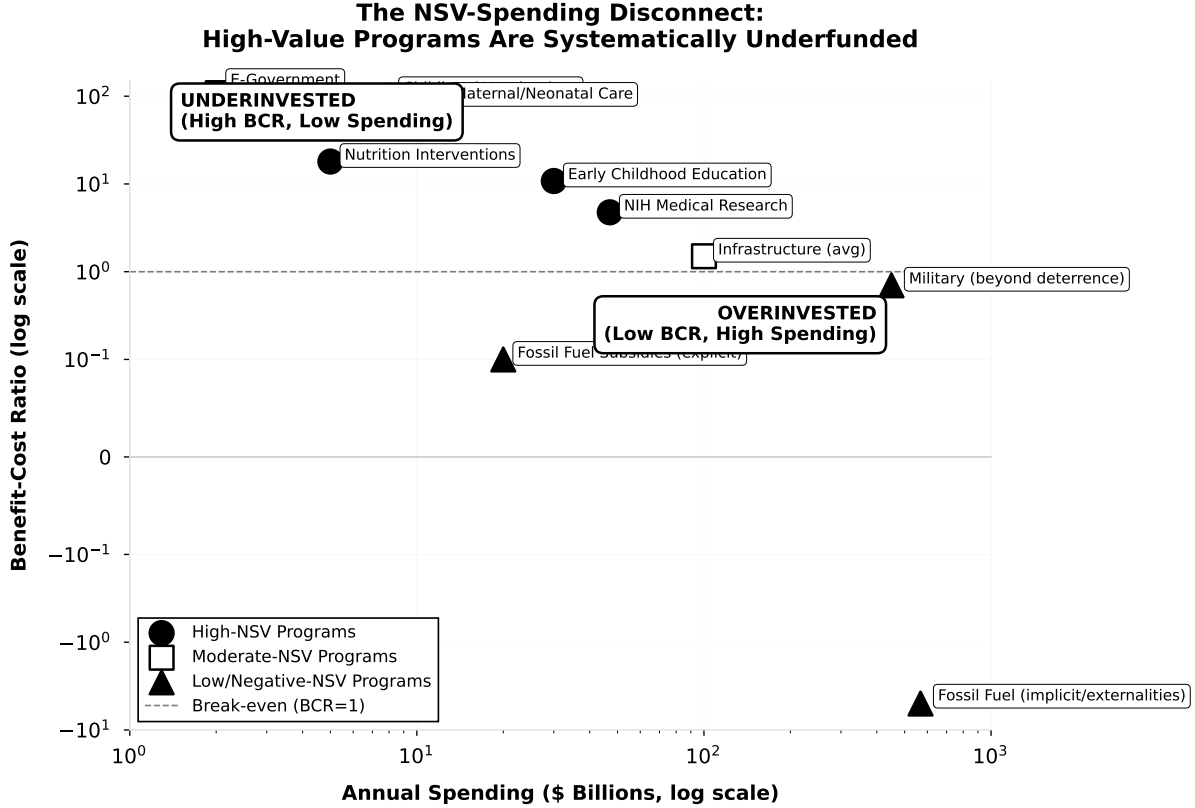


Figure 3: Government Spending vs. Benefit-Cost Ratio. Programs with the highest BCRs ($>100:1$) receive the least funding, while programs with negative NSV receive hundreds of billions. This inverse correlation demonstrates that spending is driven by lobbying intensity, not social value.

Key Insight: Under rational resource allocation, we would expect high-BCR programs to receive high spending (upper-right quadrant) and low-BCR programs to receive low spending (lower-left quadrant). Instead, we observe the opposite: programs with BCRs exceeding 100:1 are clustered in the “UNDERINVESTED” region, while programs with BCRs below 1:1 (or negative) are in the “OVERINVESTED” region.

This systematic misallocation is the empirical foundation for Incentive Alignment Bonds: **spending correlates with lobbying intensity, not social value.**

Interpretation:

- **Quadrant I (Rational):** Programs with high NSV receiving high spending: the desired equilibrium
- **Quadrant II (Underinvested):** High-NSV programs receiving low spending: Olsonian failure where diffuse benefits cannot coordinate

**The Olsonian Quadrant Chart:
Where Government Resources Flow vs. Where They Should Flow**

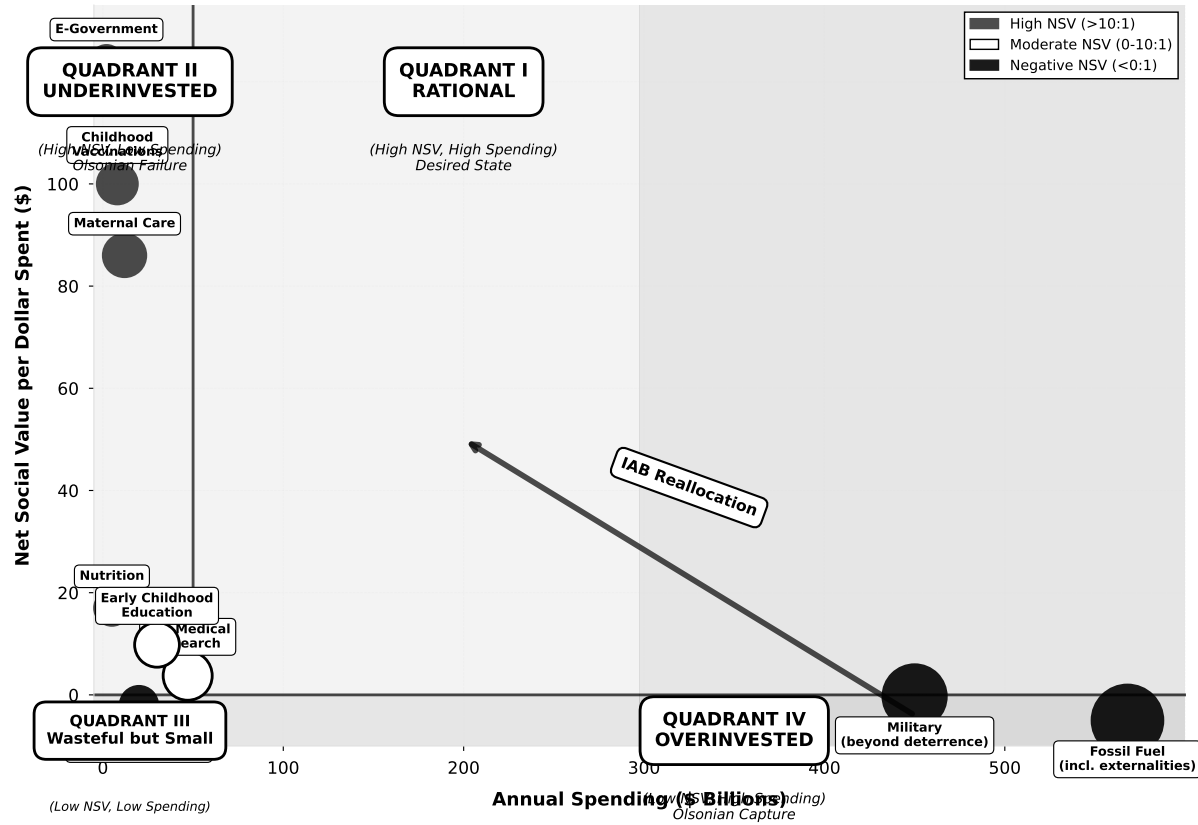


Figure 4: The Olsonian Quadrant Chart: Government Program Allocation by NSV and Spending. Quadrant I (upper-right) represents rational allocation; Quadrant II (upper-left) shows high-value programs that are systematically underfunded; Quadrant IV (lower-right) shows low-value programs that are overfunded due to Olsonian dynamics.

- **Quadrant III (Wasteful but Small):** Low-NSV programs with low spending: wasteful but limited damage
- **Quadrant IV (Overinvested):** Low-NSV programs receiving high spending: Olsonian capture where concentrated interests dominate

The IAB Mechanism: Redirects resources from Quadrant IV \rightarrow Quadrant I by making politician support for high-NSV programs individually rational through score-dependent electoral and career benefits.

Figure 3 visualizes this systematic misallocation as a scatter plot, while Figure 4 categorizes programs into four quadrants based on their NSV and funding levels. These visualizations make clear that under rational allocation, high-BCR programs receive high funding (upper-right quadrant). Instead, we observe the opposite: high-BCR programs cluster in the “UNDERINVESTED” region while low/negative-BCR programs are “OVERINVESTED.”

2.7 Relation to Existing Literature

This paper combines three literatures. First, **mechanism design theory** (Hurwicz 1972; Myerson 1979; Maskin 1999) provides tools for designing institutions when agents have private information and act strategically. While mechanism design has transformed market institutions, from spectrum auctions (Milgrom 2004) to kidney exchanges (Roth 2002), its application to political institutions is limited. We extend this framework to political incentives, treating politicians as strategic agents whose actions (votes) are observable but whose preferences may diverge from social welfare.

Second, **public choice theory** (Buchanan and Tullock 1962; Olson 1965) models politicians as utility maximizers rather than benevolent planners. Olson’s analysis of concentrated versus diffuse interests explains persistent policy failures but offers limited remedies. We turn Olson’s insight into a mechanism: if concentrated benefits cause politicians to favor narrow interests, then *concentrating benefits on the welfare-improving side* can reverse the dynamic.

Third, **campaign finance research** (Spenkuch and Toniatti 2018; Stratmann 2002; Fowler, Garro, and Spenkuch 2020) and **political agency models** (Shepherd and You 2019) document how electoral incentives and career concerns shape politician behavior. We synthesize these findings into a mechanism that systematically exploits these channels for public goods.

Relative to Social Impact Bonds (Liebman 2011; “Measuring the Success of Impact Bonds” 2017), IABs target a different agent (politicians vs. service providers) at a different level (policy adoption vs. program delivery). Relative to lobbying and campaign contributions, IABs create universal, ex-ante, public rules tied to verifiable outcomes rather than ad hoc relationships. Relative to pure voting reform proposals, IABs work within existing institutional constraints.

2.8 Contribution

This paper makes four contributions:

1. **Mechanism Design for Governance:** We apply the formal tools of mechanism design theory, developed for auctions, matching markets, and resource allocation, to the problem of political incentives. While mechanism design has transformed economic institutions from spectrum auctions to kidney exchanges, its application to democratic governance is underdeveloped.
2. **A New Financial Instrument:** We define IABs as a distinct class of financial instrument with three primitive properties: investor alignment, politician alignment, and funding from

lower-value sources. This distinguishes IABs from Social Impact Bonds (which reward service providers, not politicians), lobbying (which serves narrow interests), and campaign contributions (which lack systematic, ex-ante, universal rules tied to measurable outcomes).

3. **A Legal Architecture:** We specify a three-layer structure (scoring, electoral support, post-office benefits) that achieves incentive alignment without violating anti-bribery statutes. The legal innovation is separating data provision from electoral activity from career benefits, with each layer operated by legally distinct entities.
4. **A Welfare-Improvement Criterion:** We formalize how to identify appropriate IAB funding sources using comparative net social value analysis ($NSV_{source} < NSV_{target}$) and show why *specified reallocation*, rather than budget addition, is enough (needed) to achieve real resource shifts. Under soft budget constraints, nominal spending increases are diluted by deficit expansion and inflation while Olsonian programs retain their share of real resources. This comparative criterion implies a natural ranking of expenditures; IABs draw from the lowest-ranked programs first, ensuring reallocation is welfare-improving under standard utilitarian assumptions.

2.9 Roadmap

Section 2 reviews the relevant literature on mechanism design, public choice, Social Impact Bonds, and campaign finance. Section 3 presents the formal model, defining the politician’s utility function and proving incentive compatibility. Section 4 details the three-layer architecture and its legal basis. Section 5 analyzes potential failure modes and limitations. Section 6 discusses generalization beyond health policy to other global coordination problems. Section 7 concludes.

3 Literature Review

3.1 Mechanism Design Theory

The Sveriges Riksbank Prize in Economic Sciences 2007 was awarded to Leonid Hurwicz, Eric Maskin, and Roger Myerson “for having laid the foundations of mechanism design theory” (Committee 2007). Mechanism design addresses a fundamental problem: how should institutions be designed when agents have private information and may act strategically? Or, in plain English: how do you get selfish people to do the right thing without forcing them?

Hurwicz (1972) introduced the concept of **incentive compatibility**: a mechanism is incentive-compatible if truthful revelation (or, more generally, the socially optimal action) is in each agent’s self-interest. Translation: the rules make doing the right thing also the profitable thing. The **revelation principle** (Myerson 1979) shows that any equilibrium outcome achievable by any mechanism can be achieved by a direct mechanism in which agents truthfully reveal their private information. Maskin’s **implementation theory** (1999) clarifies when social choice rules can be implemented by mechanisms with unique equilibria.

These concepts have been applied extensively to market design: auction theory (Milgrom 2004), matching markets (Roth 2002), and regulation (Laffont and Tirole 1993). However, application to political institutions remains limited. This paper extends mechanism design to the domain of political incentives, treating politicians as strategic agents whose private information includes their true policy preferences and whose actions (votes, public statements, bill sponsorship) are observable.

3.2 Public Choice Theory

Public choice theory applies economic methods to political behavior. Buchanan and Tullock (1962) model politicians and voters as rational utility maximizers, not benevolent social planners. Olson (1965) shows that collective action for public goods is systematically undersupplied because:

1. **Free-rider problem:** Large groups cannot exclude non-contributors from benefits
2. **Asymmetric stakes:** Per-member stakes are higher in small groups
3. **Organization costs:** Smaller groups face lower coordination costs

The result is that “concentrated minor interests will be overrepresented and diffuse majority interests trumped” (Olson 1965). Empirical support includes Lohmann’s observation that U.S. sugar import quotas generated 2,261 jobs while reducing overall welfare by \$1.162 billion, an implicit cost per job exceeding \$500,000 (Lohmann 1998). The government could have paid each sugar worker half a million dollars to not grow sugar. This would have saved money. Nobody did this. Sugar growers have lobbyists. Consumers don’t.

IABs address this directly: by concentrating benefits on politicians who support diffuse-benefit policies, the asymmetry is reversed. Supporting the public good becomes the concentrated-benefit option.

3.3 Social Impact Bonds

Social Impact Bonds (SIBs), introduced in the UK in 2010, are outcome-based contracts in which private investors fund social interventions and are repaid by government only if specified outcomes are achieved (Liebman 2011). The Peterborough Prison SIB, the first implemented, funded prisoner rehabilitation and measured recidivism rates against a control group (“Lessons Learned from the Planning and Early Implementation of the Social Impact Bond at HMP Peterborough” 2011).

SIBs have attracted much policy enthusiasm but limited empirical evidence of their distinctive effect. As Hevenstone (2023) notes, “only program effects have been estimated, not the specific impact of SIB financing itself.” A systematic review by the Brookings Institution found “insufficient evidence as to whether and how SIBs deliver better outcomes than conventional forms of financing” (“Measuring the Success of Impact Bonds” 2017).

More important, SIBs target the wrong level. They incentivize **service providers** (nonprofits delivering programs) for **local outcomes** (recidivism in one city). IABs target **politicians** for **policy adoption** at the national or international level. The difference: incentivizing one job training program versus incentivizing the legislation that funds all job training programs. Table 1 summarizes the distinction.

Table 4: Comparison of Social Impact Bonds and Incentive Alignment Bonds

Dimension	Social Impact Bonds	Incentive Alignment Bonds
Target agent	Service providers	Politicians
Outcome measured	Program delivery	Policy adoption and funding flows
Scale	Municipal/program	National/international
Funding source	Government pays for outcomes	Policy outcome funds mechanism
Attribution	Single provider, single program	Voting records, funding contributions

3.4 Campaign Finance and Political Behavior

Empirical research on campaign finance provides the evidentiary foundation for the electoral layer of IABs.

Independent expenditures affect elections. Following *Citizens United v. FEC* (2010), independent expenditures increased dramatically. In the decade prior to the decision, outside groups spent \$296 million total on independent expenditures; in the decade after, they spent \$4.26 billion, a 14-fold increase (“More Money, Less Transparency: A Decade Under Citizens United” 2020). Research using transaction-level disbursement data finds that “spending on messages to voters has a statistically significant effect on voter support for candidates” and is “especially effective in changing the composition of voters” (Spenkuch and Toniatti 2018).

Interest group scorecards influence behavior. Organizations like the National Rifle Association (NRA), League of Conservation Voters (LCV), and Chamber of Commerce publish voting scorecards that affect politician behavior. The LCV has published its National Environmental Scorecard since 1970, and it has “become the gold standard of congressional vote scoring on environmental issues” (Conservation Voters 2024). Research confirms that PAC contributions correlate with votes on relevant issues, though causality is debated (Stratmann 2002; Fowler, Garro, and Spenkuch 2020).

The revolving door affects politician incentives. Research documents significant movement between government and private sector. A 2023 study found that 32% of HHS appointees exited to industry employment (Sood et al. 2023). Among four-star military officers who retired after June 2018, over 80% went to work for the defense industry as board members, advisors, executives, consultants, or lobbyists (“March of the Four-Stars: The Role of Retired Generals and Admirals in the Arms Industry” 2023). Shepherd and You (2019) find evidence that “career concerns” about post-office employment influence legislative behavior.

3.5 Credit Rating Agencies as Governance Mechanisms

Credit rating agencies (CRAs) provide a precedent for private organizations influencing sovereign policy through reputational mechanisms. Moody’s, S&P, and Fitch assign sovereign credit ratings that directly affect borrowing costs.

Barta (2024) describes CRAs as “unelected, unappointed, unaccountable profit-seeking institutions” whose power rivals the IMF or World Bank. Downgrades can trigger crises: Greece, Ireland, and Portugal all experienced accelerated debt crises following CRA downgrades to “junk” status (“Credit Rating Agencies and Sovereign Debt: Four Proposals to Support Achievement of the SDGs” 2023).

The IAB scoring system is structurally analogous to credit ratings: an independent body publishes scores based on objective criteria, and market actors (voters, donors, employers) respond to those scores. The difference is the metric: instead of fiscal sustainability, IABs measure support for measurable public goods.

4 Formal Model

For Non-Technical Readers: You Can Skip This Section

The key insight (in plain English): We create a “Public Good Score” for each politician based on their voting record. Politicians with higher scores get: 1. Campaign support from independent political groups (money for ads) 2. Better post-office job opportunities (lucrative fellowships, think tank positions) 3. Higher public status (Wikipedia article length, media coverage)

When the benefits from a higher score (campaign money + career prospects + reputation) exceed the costs (losing defense contractor donations), supporting public goods becomes the rational, self-interested choice.

You can skip to Section 4 for the practical implementation details, or read Section 6 for the governance architecture. This section provides mathematical precision for economists who need to verify the mechanism works under formal assumptions.

This section presents a formal model of political incentives under IABs. We state explicit assumptions, derive conditions for incentive compatibility (when doing the right thing is also the profitable thing), and characterize equilibria (stable outcomes). Appendix A provides detailed proofs.

4.1 Assumptions

We maintain the following assumptions throughout:

A1 (Rational Politicians). Each politician $i \in \{1, \dots, N\}$ is a rational agent maximizing expected utility U_i over career outcomes, including reelection, post-office income, and legacy.

A2 (Observable Votes). Voting records on policy class \mathcal{P} are publicly observable and verifiable. The scoring function $f : \text{VoteRecord} \rightarrow \mathbb{R}$ is common knowledge.

A3 (Credible Commitment). The IAB mechanism can credibly commit to score-dependent payoffs: independent expenditure rules $I_i(\theta_i)$ and post-office eligibility criteria $\tau(\theta_i)$ are announced ex ante and cannot be renegotiated ex post.

A4 (Funded Mechanism). The IAB is sufficiently capitalized that the payoff differentials ΔP_i , ΔY_i induced by score changes are non-negligible relative to concentrated opposition costs c_i .

A5 (Single Policy Dimension). Politicians face a binary choice $a_i \in \{0, 1\}$ on the target policy class. Extensions to multiple dimensions are discussed in Section 6.

4.2 The Politician’s Utility Function

We model a politician i as a rational agent maximizing a utility function:

$$U_i = \alpha_i \cdot P_i(\text{reelection}) + \beta_i \cdot E_i[\text{PostOfficeIncome}] + \gamma_i \cdot S_i(\text{Legacy}) \quad (3)$$

where:

- $P_i(\text{reelection})$ is the probability of winning the next election
- $E_i[\text{PostOfficeIncome}]$ is expected lifetime earnings after leaving office
- $S_i(\text{Legacy})$ is a status/legacy function (books, buildings named, Wikipedia length)

- $\alpha_i, \beta_i, \gamma_i > 0$ are weights varying by politician but assumed positive

This formulation is consistent with public choice theory’s treatment of politicians as utility maximizers (Buchanan and Tullock 1962) and with empirical research on career concerns (Shepherd and You 2019).

4.3 The Utility Function Transformation

The key mechanism design insight is that IABs transform the politician’s utility function by introducing a dependency on net social value (NSV) rankings through an intermediate score variable.

Pre-IAB utility function: Let R denote the ranking of programs by NSV (as produced by Copenhagen Consensus, GiveWell, etc.). In the status quo:

$$U_i^{\text{pre-IAB}} = \alpha_i \cdot P_i + \beta_i \cdot Y_i + \gamma_i \cdot S_i \quad (4)$$

Critically, the ranking R appears nowhere in this function. Each component is driven by lobbying intensity and concentrated interests, not social value:

- $P_i = P_i^0 + f(\text{campaign contributions}) + g(\text{attack ads avoided})$
- $Y_i = h(\text{revolving door relationships})$
- $S_i = s(\text{partisan loyalty, donor satisfaction})$

Post-IAB utility function: The IAB mechanism introduces a **Public Good Score** $\theta_i = f(R, \text{VoteRecord}_i)$ that translates the NSV ranking into politician-specific incentives. Each utility component becomes score-dependent:

$$U_i^{\text{post-IAB}} = \alpha_i \cdot P_i(\theta_i) + \beta_i \cdot Y_i(\theta_i) + \gamma_i \cdot S_i(\theta_i) \quad (5)$$

Now the ranking R is operative through θ_i . Politicians who vote to reallocate resources from low-NSV to high-NSV programs see their scores rise, which increases their reelection probability, post-office income prospects, and legacy value.

Explicit functional forms: The post-IAB components are specified as:

Reelection probability:

$$P_i(\theta_i) = P_i^0 + \delta \cdot (\theta_i - \bar{\theta}) + \epsilon_i \cdot I_i(\theta_i) \quad (6)$$

where P_i^0 is baseline reelection probability, $\bar{\theta}$ is the median score, $\delta > 0$ captures the direct electoral effect of scorecard visibility, and $I_i(\theta_i)$ is the independent expenditure function (a step function with discrete rewards):

$$I_i(\theta_i) = \begin{cases} +M & \text{if } \theta_i \geq \theta^{high} \\ 0 & \text{if } \theta^{med} \leq \theta_i < \theta^{high} \\ -M & \text{if } \theta_i < \theta^{med} \end{cases} \quad (7)$$

Post-office income:

$$Y_i(\theta_i) = Y(\tau(\theta_i)) \quad (8)$$

where $\tau : \mathbb{R} \rightarrow \{1, 2, 3\}$ maps scores to income tiers:

$$\tau(\theta_i) = \begin{cases} 1 & \text{if } \theta_i \geq 75 \quad (\text{Tier 1: \$500K+ annually}) \\ 2 & \text{if } 60 \leq \theta_i < 75 \quad (\text{Tier 2: \$200-400K annually}) \\ 3 & \text{if } \theta_i < 60 \quad (\text{Tier 3: \$150-300K annually}) \end{cases} \quad (9)$$

Legacy/status:

$$S_i(\theta_i) = S_0 + \lambda \cdot \theta_i \quad (10)$$

where S_0 is baseline status and $\lambda > 0$ captures the marginal status gain from higher scores (measured by metrics such as Wikipedia article length, think tank fellowships, speaking invitations, historical assessments).

Generated: /home/runner/work/disease-eradication-plan/disease-eradication-plan/knowledge/figures/utility-function-pre-iab.png

Figure 5

Before IABs: Politicians maximize utility based on lobbying intensity. The NSV ranking exists but doesn't affect their utility function, so they have no incentive to support high-NSV programs.

Generated: /home/runner/work/disease-eradication-plan/disease-eradication-plan/knowledge/figures/utility-function-post-iab.png

Figure 6

After IABs: The same utility function becomes score-dependent: each component (electoral, career, legacy) now depends on the politician's score, which is determined by their voting record on NSV-ranked programs. Supporting high-NSV programs becomes the rational, self-interested choice.

The transformation does not change what politicians optimize; it changes what optimizing points at. Same selfish utility maximization, radically different equilibrium outcome. This is the core mechanism design contribution: we align private incentives with public welfare not by appealing to altruism, but by redirecting self-interest. The diagrams above illustrate this transformation (compare Figure 5 with Figure 6).

4.4 Numerical Calibration and Parameter Ranges

To evaluate whether Equation 11 (the incentive compatibility condition) can realistically be satisfied, we provide illustrative calibrations based on campaign finance research and career trajectories.

Base case parameters:

Table 5: Politician Heterogeneity in Utility Weights

Politician Type	α_i (Electoral Weight)	β_i (Career Weight)	γ_i (Legacy Weight)	P_i^0 (Baseline Prob.)
Marginal seat	0.6	0.3	0.1	0.50
Safe seat	0.2	0.5	0.3	0.85
Termed out	0.0	0.7	0.3	N/A

IAB mechanism parameters (illustrative):

- Score threshold for high tier: $\theta^{high} = 80$, $\theta^{med} = 60$
- Independent expenditure support: $M = \$5\text{M}$ to $\$20\text{M}$ (competitive race)
- Electoral impact: $\delta = 0.02$ to 0.05 (2-5 percentage point swing from scorecard visibility)
- Post-office income differential: $\Delta Y = \$200\text{K}$ to $\$350\text{K}$ annually (present value $\$3\text{M}$ - $\$5\text{M}$ over 15-year career)
- Legacy multiplier: $\lambda = 0.1$ (1 SD score increase \rightarrow 10% increase in Wikipedia length, fellowships, etc.)

Concentrated opposition cost:

- Moderate opposition: $c_i = \$2\text{M}$ to $\$5\text{M}$ (attack ads, lost contributions, primary challenge risk)
- Strong opposition: $c_i = \$10\text{M}$ to $\$20\text{M}$ (defense/fossil fuel reallocation)

Incentive compatibility calculation (marginal seat politician):

Suppose a marginal-seat Senator (50% baseline reelection, $\alpha_i = 0.6$, $\beta_i = 0.3$, $\gamma_i = 0.1$) votes for a 1% military reallocation treaty:

$$\Delta P_i = \delta \cdot \Delta \theta + \epsilon_i \cdot M = (0.03)(20) + (0.08)(\$10\text{M}) = 0.6 + \$0.8\text{M} = 1.4\%$$

$$\Delta Y_i = \$300\text{K annually} \times 15 \text{ years} \times 0.05 \text{ discount} \approx \$4.5\text{M PV}$$

$$\Delta S_i = \lambda \cdot \Delta \theta = 0.1 \times 20 = 2.0$$

Utility gain:

$$\alpha_i \Delta P_i + \beta_i \Delta Y_i + \gamma_i \Delta S_i = 0.6(0.014) + 0.3(\$4.5\text{M}) + 0.1(2.0)$$

$$\approx 0.0084 + \$1.35\text{M} + 0.2 \approx \$1.35\text{M}$$

If concentrated opposition cost is $c_i = \$5\text{M}$ (attack ads, lost defense PAC contributions), then:

$$\Delta U_i = \$1.35\text{M} - \$5\text{M} = -\$3.65\text{M} < 0$$

Mechanism fails with current calibration. To achieve incentive compatibility, IABs must increase either:

- Independent expenditure support (M) from \$10M to \$20M+
- Post-office income differential (ΔY) via additional tier benefits
- Electoral impact (δ) through more aggressive scorecard visibility campaigns

Revised calibration (mechanism succeeds):

If $M = \$20M$ and $\delta = 0.05$:

$$\Delta P_i = 0.05(20) + 0.08(\$20M) = 1.0 + 1.6 = 2.6\%$$

$$\Delta U_i = 0.6(0.026)(\text{value of 2.6pp reelection}) + 0.3(\$4.5M) + 0.1(2.0)$$

Translating reelection probability to dollars (value of Senate seat = \$50M lifetime value):

$$\Delta U_i \approx 0.6(0.026 \times \$50M) + \$1.35M + 0.2 \approx \$0.78M + \$1.35M = \$2.13M$$

Still insufficient if $c_i = \$5M$. But if concentrated opposition is only $c_i = \$2M$ (lower-opposition domain like pandemic preparedness), mechanism succeeds:

$$\Delta U_i = \$2.13M - \$2M = \$0.13M > 0 \quad \checkmark$$

Key insights from calibration:

1. **Mechanism works better in lower-opposition domains:** Pandemic preparedness ($c_i = \$1M$ -\$3M) more tractable than defense reallocation ($c_i = \$10M$ -\$20M)
2. **Safe-seat politicians harder to move:** Need larger β_i (career) and γ_i (legacy) incentives since α_i (electoral) is small
3. **Termed-out politicians are most cost-effective:** $\alpha_i = 0$ but $\beta_i = 0.7$ means post-office incentives dominate; cheaper to influence
4. **Funding requirements scale with opposition:** High-opposition domains require \$20M+ per pivotal vote; low-opposition domains require \$5M-\$10M

This calibration demonstrates that Equation 16 is achievable with realistic parameter values in carefully-selected domains, validating the theoretical mechanism.

4.5 The Policy Choice

Consider a binary policy choice $a_i \in \{0, 1\}$ where:

- $a_i = 1$: Support a policy that funds public good G
- $a_i = 0$: Oppose or abstain

Let $W(G)$ denote the social welfare gain from G . By assumption, $W(G) > 0$; the policy is welfare-improving. The question is whether $a_i = 1$ is incentive-compatible.

4.6 The Pre-IAB Equilibrium

Without IABs, the politician faces:

Benefits of $a_i = 1$:

- Diffuse voter approval (small per-voter benefit, hard to attribute)
- Abstract “doing the right thing” utility (assumed small)

Costs of $a_i = 1$:

- Concentrated opposition from losers (defense contractors, pharmaceutical incumbents)
- Attack ads: “Senator voted to WEAKEN AMERICA”
- Loss of campaign contributions from concentrated interests

Formally:

$$\Delta U_i^{\text{pre-IAB}}(a_i = 1) = \epsilon - c_i \quad (11)$$

where ϵ is the small diffuse benefit and $c_i > 0$ is the net concentrated cost from losing incumbents’ support (campaign contributions foregone, attack ads received, post-office opportunities closed). Since $c_i > \epsilon$ for most policies with diffuse benefits, the equilibrium is $a_i^* = 0$. This is Olson’s result in formal terms.

4.7 The IAB Mechanism

The IAB mechanism introduces a **Public Good Score** θ_i for each politician, where:

$$\theta_i = f(\text{VoteRecord}_i) \quad (12)$$

The score is based purely on voting record on policy class \mathcal{P} (policies meeting specified criteria for the target public good). This is a design choice with important implications:

1. **Measurability:** Voting records are public, verifiable, and effectively ungameable
2. **Attribution:** Each politician’s vote is directly attributable
3. **No oracle problem:** No need for contested measurement of downstream outcomes; the mechanism relies on observable voting records

4.8 Score-Dependent Payoffs

The IAB mechanism makes each component of U_i a function of θ_i :

Reelection probability:

$$P_i(\text{reelection}) = P_i^0 + \delta \cdot (\theta_i - \bar{\theta}) + \epsilon_i \cdot I_i(\theta_i) \quad (13)$$

where:

- P_i^0 is baseline reelection probability
- $\bar{\theta}$ is the median score
- $\delta > 0$ is an empirical parameter capturing the direct electoral effect of scorecard visibility (media coverage, voter information)

- $I_i(\theta_i)$ is independent expenditure support determined by a pre-announced, public rule, a step function:

$$I_i(\theta_i) = \begin{cases} +M & \text{if } \theta_i \geq \theta^{high} \\ 0 & \text{if } \theta^{med} \leq \theta_i < \theta^{high} \\ -M & \text{if } \theta_i < \theta^{med} \end{cases} \quad (14)$$

where $M > 0$ represents campaign support magnitude.

Post-office income:

$$E_i[\text{PostOfficeIncome}] = Y(\tau(\theta_i)) \quad (15)$$

where $\tau : \mathbb{R} \rightarrow \{1, 2, 3\}$ is a tier function and $Y(1) > Y(2) > Y(3)$ represents expected annual income by tier:

Table 6: Post-Office Income by Public Good Score Tier

Tier	Threshold	Expected Annual Income	Examples
1	$\theta_i \geq 75$	\$500K+	WHO Advisory Board, Aspen Fellowships
2	$60 \leq \theta_i < 75$	\$200-400K	Brookings, RAND, university chairs
3	$\theta_i < 60$	\$150-300K	Defense contractor boards, lobbying firms

4.9 Incentive Compatibility

Proposition 1 (Sufficient Condition for Incentive Compatibility). *Under assumptions A1–A5, if the score gain from supporting policy class \mathcal{P} is $\Delta\theta > 0$, and*

$$\alpha_i \cdot \Delta P_i + \beta_i \cdot \Delta Y_i + \gamma_i \cdot \Delta S_i > c_i \quad (16)$$

then $a_i = 1$ is the unique best response for politician i .

Proof sketch: By A2, votes are observable and the scoring function is common knowledge, so the politician can compute $\theta'_i = \theta_i + \Delta\theta$ conditional on $a_i = 1$. By A3, the payoff functions are credibly committed, so the politician can compute ΔP_i , ΔY_i , and ΔS_i . By A1, the politician maximizes U_i . The change in utility from choosing $a_i = 1$ versus $a_i = 0$ is:

$$\Delta U_i = \alpha_i \cdot \Delta P_i + \beta_i \cdot \Delta Y_i + \gamma_i \cdot \Delta S_i - c_i$$

When Equation 16 holds, $\Delta U_i > 0$, so $a_i = 1$ strictly dominates $a_i = 0$. See Appendix A for the complete proof. \square

Corollary 1. *Under A4, there exists a funding level \bar{F} such that for all $F > \bar{F}$, Equation 16 holds for all politicians with $c_i < \bar{c}$ for some threshold $\bar{c}(F)$ increasing in F .*

This establishes that sufficiently funded IABs can overcome concentrated opposition up to a threshold that increases with funding.

4.10 Nash Equilibrium Analysis

Consider a legislature of N politicians. Let $n = \sum_{i=1}^N a_i$ be the number supporting the policy.

Proposition 2 (Multiple Equilibria Without IABs). *Under A1–A2 and A5, without the IAB mechanism, the game among N politicians has at least two pure strategy Nash equilibria:*

- (i) *The all-defect equilibrium $(a_1, \dots, a_N) = (0, \dots, 0)$*
- (ii) *Potentially the all-cooperate equilibrium $(a_1, \dots, a_N) = (1, \dots, 1)$ if coordination is feasible*

The all-defect equilibrium is risk-dominant when $c_i > \epsilon$ for all i .

Proof sketch: In the all-defect equilibrium, no politician benefits from unilateral deviation because the diffuse benefit ϵ is outweighed by the concentrated cost c_i . The all-cooperate equilibrium may exist if coordination reduces per-politician costs or if ϵ aggregates across politicians, but it is unstable to individual defection when concentrated interests can target defectors. See Appendix A. \square

Proposition 3 (Equilibrium Selection With IABs). *Under A1–A5, if the IAB mechanism is funded such that Equation 16 holds for all i , then $(1, \dots, 1)$ is the unique Nash equilibrium.*

Proof sketch: When Equation 16 holds for each politician i , choosing $a_i = 1$ is a strictly dominant strategy regardless of other politicians’ choices. A profile of strictly dominant strategies constitutes the unique Nash equilibrium. \square

Remark. The “dominant strategy” characterization applies to the stylized binary choice taking the IAB mechanism as given. In richer settings with endogenous IAB design, strategic scoring manipulation, or multiple policy dimensions, additional equilibrium refinements apply. Section 5 discusses these extensions.

4.11 Illustrative Example: A Global Health Treaty

To ground the formal model in concrete terms, consider a hypothetical application. Suppose one’s public goods objective is to simultaneously reduce global conflict and reduce the global burden of disease. Global military expenditure currently exceeds \$2.4 trillion annually (“Trends in World Military Expenditure, 2023” 2024), while the global burden of disease exceeds 2.5 billion disability-adjusted life years (DALYs) annually (“Global Health Estimates: Life Expectancy and Leading Causes of Death and Disability” 2024). An international treaty in which signatory nations commit to redirecting 1% of military spending to a global pragmatic clinical trial system generates approximately \$24 billion per year for pragmatic clinical trials, roughly tripling current global clinical trial funding, while creating modest but meaningful pressure toward demilitarization.

The IAB mechanism for this treaty would allocate treaty inflows as follows:

- **80%** to the public good itself (pragmatic clinical trials)
- **10%** to investor returns (perpetual payments to those who funded the campaign to pass the treaty)
- **10%** to political incentives (funding the three-layer architecture)

This allocation structure creates aligned incentives across all participants. Investors who funded the initial campaign receive perpetual returns (\$2.4B annually) as long as the treaty continues, giving them strong incentives to support treaty expansion and defend against repeal. The political incentive allocation (\$2.4B annually) funds the three-layer architecture:

Scoring layer: Politicians receive Health Research Scores based on their voting record on treaty ratification, annual funding appropriations, and related legislation. A legislator who votes YES on the treaty and subsequent funding bills sees their score rise; one who votes NO sees their score fall.

Electoral layer: With \$2.4 billion annually available, independent expenditure campaigns can credibly commit: “We will spend \$50 million supporting high-scorers in competitive races.” At typical costs per competitive race, this funds meaningful independent campaigns in 20-30 races annually.

Post-office layer: Foundations funded by the political incentive allocation establish eligibility criteria: “The Global Health Leadership Fellowship (\$300K/year, 5-year term) requires a career Health Research Score above 70.”

The self-funding nature matters. The treaty’s success generates the resources that sustain both investor returns and political incentives for its continuation and expansion. This creates multiple reinforcing feedback loops:

1. **Investor pressure for expansion:** Investors receiving 10% of \$24B want 10% of \$48B (if the treaty expands to 2%), creating a constituency that lobbies for treaty growth
2. **Political incentives for continuation:** Politicians who supported the treaty benefit from ongoing electoral and career support, incentivizing them to defend it against repeal
3. **Escalation dynamics:** Each expansion (1% \rightarrow 2% \rightarrow 5%) increases both investor returns and political incentive funding, strengthening the coalition for further expansion

In other words, we’re creating a lobbying machine for public goods that gets stronger the more successful it becomes. This is the same dynamic that made defense contractors powerful, except pointed at curing diseases instead of building bombs. Once the flywheel starts spinning, concentrated opposition faces an opponent that grows with every victory.

Consider the decision calculus for a hypothetical Senator Smith facing a vote on treaty ratification:

Table 7: Senator Smith’s Decision Calculus

Without IABs	With IABs
Defense contractors fund opponent if YES	Health Research Score rises 25 points
Attack ads: “Smith weakened our military”	Independent campaigns spend \$2M supporting Smith
Benefits (cures) arrive in 10+ years	Post-office eligibility: Tier 3 \rightarrow Tier 1
Diffuse beneficiaries cannot coordinate	Expected post-office income: \$200K \rightarrow \$400K/yr

The IAB mechanism transforms the incentive landscape. The concentrated costs (defense contractor opposition) remain, but they are now outweighed by concentrated benefits (score increases, electoral support, career advancement). Supporting the treaty becomes the utility-maximizing choice.

4.12 Calibration: Parameter Ranges for Incentive Compatibility

To assess whether IABs can achieve incentive compatibility, we calibrate the model using empirical estimates from campaign finance and lobbying research. Table 8 reports parameter ranges.

Table 8: Calibration Parameters for IAB Model

Parameter	Symbol	Empirical Range	Source
Defense contractor opposition spending	c_i	\$0.5–5M per race	OpenSecrets 2020
Independent expenditure effect on vote share	ϵ_i	0.5–2 pp per \$1M	Spenkuch & Toniatti 2018
Value of 1 pp reelection probability	α_i	\$0.5–2M	Implied by campaign spending
Post-office income differential (Tier 1 vs 3)	ΔY	\$150–300K/yr	Industry salary data
Discount rate for post-office income	r	5–10%	Standard
Career length post-office	T	10–20 years	Empirical average

Calibration exercise. Consider a senator facing concentrated opposition of $c_i = \$2\text{M}$ from defense contractors. Under the IAB mechanism with political incentive funding of \$2.4B annually:

Electoral layer: If \$50M is allocated to 25 competitive races, each high-scorer receives \$2M in independent expenditure support. At 1 pp per \$1M, this implies $\Delta P_i \approx 2$ pp. With $\alpha_i = \$1\text{M}$ per pp, the electoral benefit is $\alpha_i \cdot \Delta P_i = \2M .

Post-office layer: Moving from Tier 3 to Tier 1 increases expected annual post-office income by \$200K. Given $T = 15$ years and $r = 7\%$, we calculate the present value:

$$\Delta Y_i = \$200\text{K} \times \frac{1 - (1.07)^{-15}}{0.07} \approx \$1.8\text{M}$$

Total benefit: $\alpha_i \cdot \Delta P_i + \beta_i \cdot \Delta Y_i \approx \$2\text{M} + \$1.8\text{M} = \3.8M

Incentive compatibility check: Since $\$3.8\text{M} > c_i = \2M , the incentive compatibility condition Equation 16 is satisfied. The senator strictly prefers supporting the treaty.

This calibration suggests that IAB funding levels in the billions can overcome concentrated opposition in the millions, consistent with the “lobbying alpha” asymmetry observed empirically.

5 The Three-Layer Architecture

To implement score-dependent payoffs, we need mechanisms. This section specifies a three-layer architecture and analyzes its legality under U.S. law.

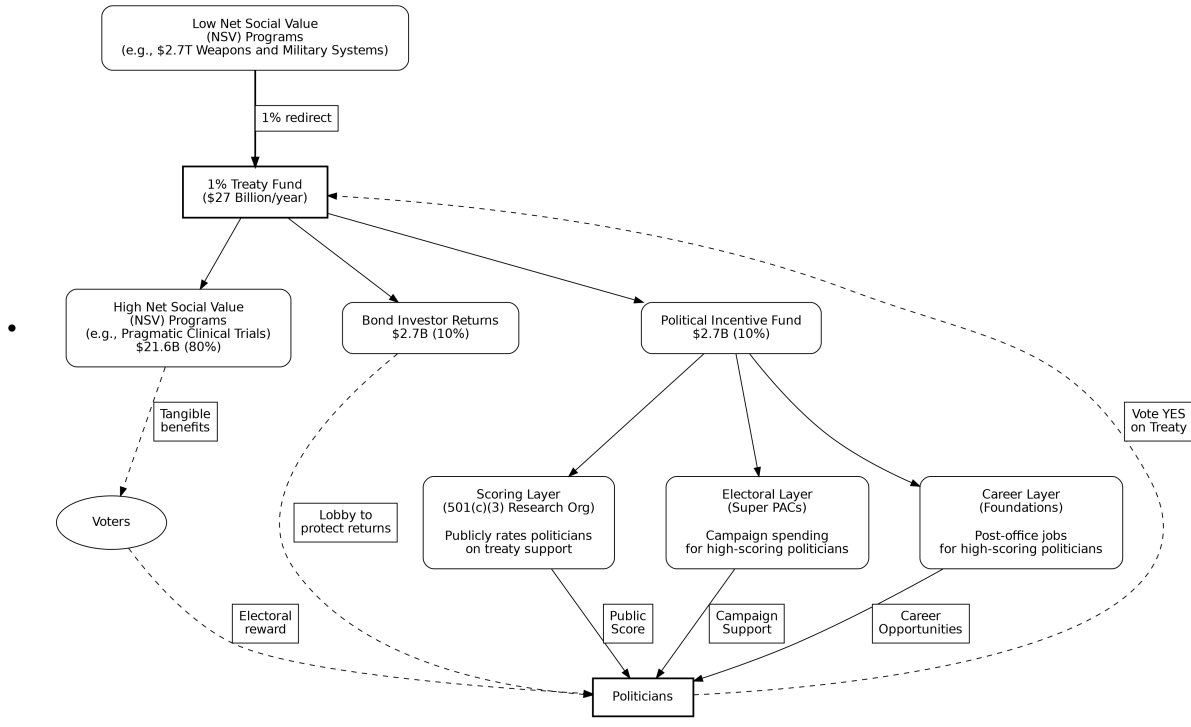


Figure 7

5.1 Layer A: Scoring (Data Provision)

An independent entity, structured as a 501(c)(3) research organization, maintains the Public Good Score system.

Inputs:

- Public voting records from official government sources
- Bill classifications (does legislation meet criteria for policy class \mathcal{P} ?)
- Sponsorship and co-sponsorship records

Outputs:

- Individual politician scores θ_i
- Historical score trajectories
- Comparative rankings

Legal basis: 501(c)(3) organizations are permitted to conduct research and publish findings. Publishing voting records and scores based on those records is protected speech. The key constraint: 501(c)(3) organizations cannot “participate in, or intervene in any political campaign on behalf of (or in opposition to) any candidate for public office” (Service).

The scoring layer does not violate this constraint because it: 1. Applies universal, pre-announced criteria 2. Does not endorse or oppose specific candidates 3. Reports data that is already public

5.2 Layer B: Electoral Support

Independent actors, structured as 501(c)(4) organizations, PACs, or Super PACs, commit to electoral support rules tied to scores.

Mechanism:

- Super PACs announce: “We will spend \$X million on independent expenditures supporting candidates with Public Good Score ≥ 70 in competitive races”
- This is a standing, public, ex-ante rule, not a deal with any specific candidate

Legal basis: *Citizens United v. FEC* (2010) held that independent expenditures are protected speech and cannot be limited (United States 2010). The Court defined corruption narrowly as “quid pro quo,” meaning an explicit exchange of money for official acts. Independent expenditures, not coordinated with candidates, cannot constitute quid pro quo corruption under this standard.

Justice Kennedy wrote that independent spending “does not give rise to corruption or the appearance of corruption” because it is “not prearranged and coordinated” with candidates (United States 2010). The IAB mechanism satisfies this requirement: support is determined by a public score, not by private arrangements.

5.3 Layer C: Post-Office Benefits

Private foundations, think tanks, and advisory boards establish eligibility criteria tied to scores.

Mechanism:

- Foundations announce: “Eligibility for the Global Health Fellowship requires a career average Public Good Score ≥ 75”
- Think tanks announce: “Advisory board positions are restricted to former officials from jurisdictions that achieved threshold improvements under their leadership”

Legal basis: Private organizations have broad discretion in setting employment and fellowship criteria. Setting criteria based on public, objective metrics is standard practice. The constraint: criteria must be announced in advance and apply universally, not targeted at specific individuals post hoc.

5.4 Why This Is Not Bribery

The legal analysis proceeds through three levels: statutory elements, constitutional protections, and established precedent.

5.4.1 Statutory Analysis: 18 U.S.C. § 201

Under 18 U.S.C. § 201, bribery requires:

1. A **thing of value** given to a public official
2. **With intent to influence** an official act
3. As a **quid pro quo**, a specific exchange (Code)

The IAB mechanism fails each element:

Element 1: No thing of value to officials.

- *Scores* are information, not things of value. Publishing data about voting records is protected speech.
- *Electoral support* goes to campaigns, not officials personally. Independent expenditures benefit candidates electorally but are not personal enrichment.
- *Post-office opportunities* are future employment contingent on (a) leaving office and (b) meeting publicly-announced, universally-applied criteria. Future employment prospects are not “things of value” under bribery law. Otherwise, any industry that hires former regulators would be guilty of bribery.

Element 2: No specific intent to influence.

Rules are universal and ex-ante. No one tells any official “vote for X and receive Y.” The rule is: “Any official who supports this class of policies will score higher.” This is categorically different from “I will give you money if you vote yes on H.R. 1234.”

The distinction matters legally. In *United States v. Sun-Diamond Growers* (1999), the Supreme Court held that illegal gratuities require a link to a “specific official act,” not merely a general desire to curry favor. IABs create no such link; they reward policy *classes*, not specific votes.

Element 3: No quid pro quo.

No exchange occurs with any official. Officials are not party to any agreement. They merely observe standing rules and act in their self-interest. The Supreme Court in *McCutcheon v. FEC* (2014) held that “ingratiation and access... are not corruption” and that only “quid pro quo corruption” justifies regulation of political speech.

McDonnell v. United States (2016) further narrowed the definition of “official act,” requiring a formal exercise of governmental power on a specific matter (United States 2016). The IAB mechanism does not pay for specific official acts; it creates a scoring system that independent actors may choose to use.

5.4.2 Constitutional Protections

The IAB architecture is protected by the First Amendment at multiple levels:

1. **Scoring layer:** Publishing voting records and scores based on those records is core political speech. The government cannot prohibit citizens from evaluating and publicizing legislators’ voting records.
2. **Electoral layer:** *Citizens United* held that independent expenditures are protected speech. The government cannot limit spending by independent actors to support or oppose candidates based on their policy positions.
3. **Post-office layer:** Private organizations have First Amendment associational rights to set their own membership and employment criteria. A foundation requiring fellows to have supported certain policies is no different from a think tank preferring scholars who share its intellectual orientation.

5.4.3 Established Precedent: Decades of Legal Scorecard Operations

The most powerful evidence that IABs are legal is that their core components have operated openly for decades without prosecution:

Scoring systems:

- The **League of Conservation Voters** has published its National Environmental Scorecard since 1970, for 54 years rating legislators on environmental votes
- The **NRA** grades legislators A through F on gun rights votes and publicizes the ratings
- The **Chamber of Commerce**, **AFL-CIO**, **ACLU**, and dozens of other organizations publish voting scorecards

Electoral support tied to scores:

- The NRA explicitly endorses candidates based on their grades and spends millions on independent expenditures supporting high-scorers
- The LCV endorses based on scorecard performance
- Labor unions support candidates with pro-labor voting records

Post-office benefits tied to policy positions:

- Defense contractors hire former Pentagon officials who supported procurement programs
- Pharmaceutical companies hire former FDA officials who approved their drugs
- Think tanks hire former legislators who championed their policy priorities

None of these activities have been prosecuted as bribery. The IAB architecture merely *systematizes* what already happens ad hoc, making the rules transparent, universal, and tied to measurable public goods rather than narrow interests.

5.4.4 The Key Legal Distinction

Bribery corrupts official judgment by introducing private benefit that conflicts with public duty. IABs align private benefit with public duty, rewarding officials for *doing their job well* as measured by policy outcomes. This is the opposite of corruption.

The legal architecture ensures this alignment holds:

- No money flows to officials while in office
- Rules apply universally to all officials ex ante
- Criteria are based on public records, not private arrangements
- Benefits come from independent actors, not parties to any agreement with officials

5.5 Legal Entity Separation

Table 9: Legal Entity Separation in the IAB Architecture

Layer	Entity Type	Permitted Activities	Prohibited Activities
Scoring	501(c)(3)	Research, publish scores	Campaign intervention
Electoral	501(c)(4), PAC, Super PAC	Independent expenditures	Coordination with candidates
Post-office	Private foundations	Set eligibility criteria	Condition grants on specific votes

5.6 Funding Sources and Foundation Investment

IABs can attract capital from multiple sources with different legal constraints.

Commercial investors (impact funds, family offices, institutional capital) face no restrictions on which layers they fund. If the bond offers market-rate or above-market returns, commercial capital can fund all activities, including the electoral layer.

Private foundations seeking to invest via Program-Related Investments (PRIs) face an additional constraint. IRS regulations require that PRIs not be used “directly or indirectly to lobby or for political purposes” (Service). This means:

Layer	Foundation PRI Eligible?	Notes
Scoring	Yes	Pure research/data
Electoral	No	Explicitly political
Post-office	Yes	Employment criteria

This constraint is less significant than it appears. In a unified bond offering, the allocation happens *after* policy success. Investors receive returns from policy revenue, which is then allocated across uses. Foundations investing via PRI could have their returns earmarked for non-electoral uses.

Precedent: The Rockefeller Foundation invested in the Peterborough Social Impact Bond via PRI. The Kresge Foundation and Living Cities provided junior tranches in Massachusetts Pay for Success projects. Bloomberg Philanthropies guaranteed Goldman Sachs’ investment in the Rikers Island SIB. Foundation participation in outcome-based bonds is established practice.

Practical implication: For IABs offering strong commercial returns, foundation PRIs are a supplementary funding source, not a necessity. The electoral layer, the one foundations cannot fund, can be capitalized by return-seeking investors.

6 Failure Modes and Limitations

6.1 Gaming and Metric Corruption

Any metric can be gamed. Goodhart’s Law states: “When a measure becomes a target, it ceases to be a good measure” (Goodhart 1984). Potential gaming strategies:

1. **Cheap talk:** Politicians vote symbolically for popular positions but undermine implementation
2. **Bill stuffing:** Attach poison pills to supported legislation
3. **Strategic timing:** Time votes when outcomes are predetermined

Countermeasures:

- Score multiple dimensions (votes, sponsorship, floor statements, implementation oversight)
- Weight by vote significance (close votes count more)
- Track long-term outcomes as validation (though not as scoring basis)

6.2 Plutocracy Objection and the “Lobbying With Extra Steps” Critique

The most serious criticism economists are likely to raise: **“This is just lobbying with extra steps. You’re creating a well-funded interest group to lobby politicians for your preferred policies.”**

This criticism deserves a careful response, as it goes to the heart of whether IABs represent genuine institutional innovation or merely repackage existing capture dynamics.

What makes IABs structurally different from corporate lobbying:

1. **Comparative welfare criterion:** IABs fund reallocation from $NSV_{source} < NSV_{target}$ programs. Corporate lobbying seeks absolute budget increases for the lobbying firm's industry, regardless of comparative value. Defense contractors lobby for *more military spending* without specifying what gets cut. IABs specify *both* the source (lower-NSV) and target (higher-NSV), creating a welfare-improving constraint.
2. **Public good vs. private good alignment:** Corporate lobbying aligns politician incentives with excludable private benefits (defense contracts flow to specific firms). IAB-supported policies produce non-excludable public goods (cures, climate stability). Investors cannot capture the primary benefits. They accrue to the general public. Investors capture only a fraction of the funding flow, not the end-state benefit.
3. **Transparent, universal, ex-ante rules:** Corporate lobbying operates through opaque relationships (revolving door promises, implicit quid pro quos, insider access). IABs publish scoring criteria ex-ante, apply them universally to all politicians regardless of relationship, and make scores publicly available. Any politician can improve their score through observable actions.
4. **Diffuse vs. concentrated funders:** Corporate lobbying concentrates returns to a small number of firms. IABs, if structured as retail-accessible securities, allow millions of diffuse beneficiaries to invest small amounts. A \$1,000 IAB investment makes a retiree a “special interest” in pragmatic clinical trials, structurally impossible with current lobbying.
5. **Metric validation:** Corporate lobbying success is measured by dollars flowing to the lobbying firm. IAB success is measured by rigorous external benefit-cost analyses (Copenhagen Consensus, GiveWell, academic literature). The scoring organizations are independent 501(c)(3) entities, not the investors themselves.

However, the objection retains force: IABs do not democratize power; they redirect it. Whose conception of “public good” defines the NSV ranking? What prevents IAB funders from capturing the scoring process? What prevents redirected resources from being recaptured by new concentrated interests?

The structural answer: Wishocracy (Layer 2) solves the allocation capture problem.

As detailed in Section 6.2.3, once IABs successfully redirect resources to high-NSV domains, **Wishocracy prevents those resources from being captured** by replacing committee allocation with democratic preference aggregation:

1. **Who decides priorities?** All citizens, through pairwise comparisons aggregated via Bradley-Terry model
2. **What prevents capture of the allocation process?** Random pairwise sampling makes advertising economically infeasible ($\sim 0.1\%$ appearance probability per campaign \rightarrow \$200M advertising cost for rankings shift \rightarrow cheaper to just do the research)
3. **What prevents IAB funders from controlling the scoring process?** Layer 0 (Wishocracy for Domain Ranking) combines expert BCR data from multiple independent sources (Copenhagen Consensus, GiveWell, IMF, academic literature) with democratic aggregation via citizen pairwise comparisons. IAB funders cannot capture the domain rankings because that would require capturing all expert organizations AND manipulating millions of citizen

preferences. Layer 2 (Wishocracy) then democratizes allocation *within* domains approved via Layer 0 rankings.

The division of labor:

- **Layer 0 (Wishocracy for Domain Ranking):** Determine *which domains* should be funded (pragmatic clinical trials vs. military spending vs. climate vs. subsidies). Expert organizations provide BCR data; citizens aggregate via pairwise comparisons to produce democratic domain rankings.
- **Layer 1 (IABs):** Make politicians care about Layer 0 rankings via score-dependent incentives
- **Layer 2 (Wishocracy for Within-Domain Allocation):** Citizens allocate *within approved domains* (30% cancer, 20% Alzheimer’s, etc.) without committee capture
- **Layer 3 (Domain marketplaces):** Market mechanisms determine *specific projects* (which trials, which researchers)

Residual governance requirements for Layer 0 (Wishocracy for Domain Ranking):

- **Multiple independent expert organizations:** Copenhagen Consensus, GiveWell, IMF, CBO, academic literature provide BCR data (no monopoly on expertise)
- **Academic peer review:** BCR methodologies published in peer-reviewed journals, subject to replication
- **Transparent methodology:** All expert data sources, BCR calculations, and Wishocracy aggregation algorithms publicly available
- **Democratic aggregation:** Bradley-Terry model aggregates citizen pairwise comparisons; prevents capture of domain prioritization
- **Anti-gaming:** Random pairwise sampling makes advertising economically infeasible (same mechanism as Layer 2)

The key innovation: We don’t need a single authoritative answer to “whose conception of public good governs allocation” at Layer 1 because **Layers 0 and 2 democratize that decision**. Layer 0 democratizes *domain* prioritization (medical research vs. military spending). Layer 2 democratizes *within-domain* allocation (Alzheimer’s vs. cancer). IABs (Layer 1) solve the *adoption* problem, getting politicians to vote for reallocation based on democratically-determined priorities. This is why the four-layer stack is needed. Each layer addresses a distinct failure mode.

Democratic legitimacy: IABs overcome the Olsonian collective action failure at the *adoption* stage, the binding constraint that prevents welfare-improving policies from being considered. Democratic legitimacy is addressed at two levels: **(1) Layer 0** democratizes domain prioritization (which broad areas like medical research, climate, pandemic preparedness should be funded), and **(2) Layer 2** democratizes within-domain allocation (how to allocate within medical research across cancer, Alzheimer’s, etc.). Both layers use Wishocracy’s pairwise comparison mechanism, ensuring that resource allocation reflects aggregated citizen preferences rather than committee capture or funder control. The mechanism is compatible with democratic governance precisely because domain prioritization, adoption, and allocation are separated across layers.

6.3 Unintended Consequences

Political systems are complex. Possible unintended consequences:

1. **Crowding out intrinsic motivation:** If politicians come to see policy support as instrumental, they may become more transactional overall. (One worries that politicians might start

caring more about campaign donations and career prospects than about serving the public good. This would represent a dramatic departure from the current system, where politicians are famously motivated solely by civic duty.)

2. **Metric fixation:** Excessive focus on scored policies at the expense of unscored but important issues
3. **Legitimacy erosion:** Public perception that politicians are “bought” (even legally) may reduce trust

Given these risks, start with a careful pilot. Measure what happens. Scale if it works.

7 IABs in the Broader Governance Architecture

7.1 The Four-Layer Governance Stack

IABs do not solve all governance problems. They address one specific failure mode: **how to get welfare-improving policies adopted when concentrated interests oppose them**. To understand IABs’ role in the broader challenge of optimal resource allocation, we propose a four-layer governance stack where each layer addresses a distinct problem.

Layer 0: Wishocracy for Democratic Domain Ranking

Problem addressed: What should we fund? Whose conception of “public good” governs resource allocation?

Solution: Expert organizations provide benefit-cost data; citizens aggregate via pairwise comparisons to produce democratic domain rankings.

How it works:

1. **Expert organizations produce BCR analyses** (this already exists):

Organization	Focus	Key Outputs
Copenhagen Consensus	Development interventions	BCR rankings since 2004
GiveWell	Global health	Cost-effectiveness analyses
IMF	Subsidies and fiscal policy	Fossil fuel subsidy database
United for Medical Research	Medical research ROI	NIH impact analyses
Brown University Costs of War	Military spending	Job creation per dollar by sector
Congressional Budget Office	Federal programs	Program evaluations
PNAS / Academic Literature	Externalities	Climate, pollution valuations

2. **Citizens review data and make pairwise comparisons:** Wishocracy platform shows random pairs of domains (Pragmatic Clinical Trials vs. Military Spending). For each pair, users see synthesized data from multiple expert sources (Copenhagen Consensus BCR: 100:1 vs. 0.7:1; GiveWell analyses; IMF externality estimates; etc.). Citizens choose which domain should receive MORE resources.
3. **Bradley-Terry aggregation produces democratic rankings:** Millions of pairwise comparisons aggregate into a global ranking of domains by democratically-determined priority.

Key advantages over expert-only rankings:

- **Prevents capture:** Would need to capture ALL expert organizations + manipulate millions of citizen preferences (much harder than capturing a single scoring organization)
- **Democratic legitimacy:** Rankings reflect aggregated citizen values after reviewing expert analyses, not unelected technocrats deciding priorities
- **Value pluralism:** Citizens can weight domains based on diverse values beyond utilitarian BCR (some prioritize climate, others health, etc.)
- **Anti-gaming:** Random pairwise sampling makes advertising economically infeasible (same mechanism as Layer 2)

The original problem remains: Expert rankings exist but don't influence allocation because they don't appear in politician utility functions. **The marginal value of producing another expert ranking is zero.** Wishocracy at Layer 0 solves democratic legitimacy and capture-resistance, but doesn't solve the adoption problem. That requires Layer 1.

Layer 1: Incentive Alignment Bonds

Problem addressed: How do we make NSV rankings consequential?

Solution: Insert rankings into politician utility functions via score-dependent electoral support and career benefits.

Key innovation: IABs are the **binding constraint** in the governance stack. Layer 0 (rankings) exists but is ineffective without Layer 1. Layers 2 and 3 (below) cannot function until Layer 1 creates the political conditions for their adoption.

This is what the current paper specifies: The three-layer legal architecture (scoring, electoral, post-office) that implements this mechanism.

Layer 2: Wishocracy (Post-Adoption Allocation)

Problem addressed: Once resources are redirected to high-NSV domains, how do we allocate within those domains without recreating Olsonian capture?

Solution: Aggregated Pairwise Preference Allocation, where citizens make simple pairwise comparisons ("Curing Alzheimer's vs. Climate Change?") and a Bradley-Terry model aggregates millions of comparisons into global rankings (Wishonia 2024).

Why this is separate from IABs:

- IABs get the treaty passed and resources redirected
- Wishocracy prevents those resources from being captured by new concentrated interests
- Example: IAB treaty redirects \$24B to pragmatic clinical trials. Wishocracy determines allocation across cancer, Alzheimer's, rare diseases, etc.

Division of labor:

- **Layer 1 (IABs):** Politicians → Vote for treaty
- **Layer 2 (Wishocracy):** Citizens → Allocate across diseases/interventions

Crowdfunding Campaign Evolution:

The Wishocracy mechanism can be extended beyond government program allocation to replace traditional research grants and contracts entirely. Instead of government bureaucracies deciding which projects to fund, researchers create crowdfunding campaigns. These campaigns explain:

- Why their project creates the greatest marginal net social value per dollar
- Evidence they can implement (past results, team credentials, methodology)
- Expected outcomes and benefit-cost ratios

Citizens then allocate between these campaigns through pairwise comparisons, with the Bradley-Terry model aggregating preferences. Funding flows directly based on aggregated citizen rankings, eliminating government gatekeepers.

Anti-Gaming Through Random Pairwise Sampling:

This mechanism addresses the critical failure mode of traditional crowdfunding platforms (Kickstarter, GoFundMe, Indiegogo): **marketing skill matters more than merit**. Campaigns with professional videos, celebrity endorsements, and advertising budgets systematically outperform higher-value projects with poor presentation.

Wishocracy’s random pairwise sampling makes advertising economically infeasible.

The math: With 10,000 active campaigns and each user completing 50 pairwise comparisons, the probability any specific campaign appears in a given comparison is $\frac{2}{10,000} \approx 0.02\%$, and the probability a user sees a specific campaign in their 50 comparisons is $\approx 0.1\%$. To shift a campaign from the 15th percentile to the top 5% requires approximately 10 million favorable comparisons, which means persuading approximately 200 million users (at 50 comparisons each), at an advertising cost of approximately \$200 million to reach that many users with persuasive content.

Economic conclusion: Spending \$200 million on advertising to win research funding is irrational. **It’s cheaper to just do the research yourself.** The only cost-effective strategy is to create genuinely high-NSV projects that win pairwise comparisons on merit.

Contrast with traditional crowdfunding:

- **Kickstarter:** All campaigns visible on browse pages, search results. A \$100K ad campaign can reach millions of potential backers directly.
- **Wishocracy:** Each campaign appears in $\sim 0.1\%$ of comparisons. The same \$100K ad campaign increases funding probability by $< 0.01\%$.

Result: Only merit-based quality matters. Projects rise or fall based on their actual NSV per dollar, not their promotional budget. This solves the “allocation capture problem” at Layer 2: even after IABs redirect resources to high-NSV domains, those resources flow to the highest-value specific projects within each domain.

Layer 3: Domain-Specific Marketplaces

Problem addressed: Within a specific domain (e.g., cancer research), how do we allocate to specific projects, researchers, and trials?

Solution: Market mechanisms like a decentralized framework for drug assessment for clinical trials, prize markets for research breakthroughs, retroactive public goods funding.

Example cascade:

1. **Layer 0 (Wishocracy for Domain Ranking):** Citizens review expert BCR data (Copenhagen Consensus: pragmatic clinical trials 100:1, military 0.7:1; GiveWell analyses; IMF externality data) and aggregate via pairwise comparisons to rank pragmatic clinical trials above military spending

2. **Layer 1 (IABs):** Politicians pass treaty redirecting \$24B to pragmatic clinical trials based on Layer 0 democratic domain ranking
3. **Layer 2 (Wishocracy for Within-Domain Allocation):** Citizens allocate within pragmatic clinical trials: 30% to cancer, 20% to Alzheimer’s, 15% to infectious disease, etc.
4. **Layer 3 (a decentralized framework for drug assessment, or dFDA):** Researchers compete in prediction markets; trials funded based on predicted effectiveness

7.2 Why This Sequencing Matters

The four layers must be implemented in order:

1. **Layer 0 (Wishocracy for Domain Ranking) can be built now** using existing expert BCR data. The technology exists (Bradley-Terry aggregation) and the data exists (Copenhagen Consensus, GiveWell, IMF), but the resulting rankings are being ignored
2. **Layer 1 (IABs) is the binding constraint;** without it, nothing else happens. Politicians ignore Layer 0 rankings and won’t adopt Layers 2-3
3. **Layer 2 and 3 become possible** only after Layer 1 creates political support for new allocation mechanisms

Common mistake: Advocates focus on Layers 2 or 3 (better voting systems, prediction markets, AI allocation) without recognizing that **Layer 1 is the bottleneck**. You cannot implement Wishocracy or domain marketplaces if politicians won’t vote to fund them in the first place.

IABs solve the meta-problem: How do we bootstrap from the current Olsonian equilibrium to one where better governance mechanisms can be adopted?

7.3 Comparison to Alternative Governance Mechanisms

Table 12 compares IABs to alternative governance reform approaches. Comparison criteria: feasibility assessments, implementation timelines, capital requirements, and structural barriers. The analysis draws on historical precedents for institutional reform, capital requirements relative to available funding sources, political economy barriers, and adoption rates of analogous mechanisms.

Table 12: Comparative Feasibility Assessment of Governance Reform Mechanisms

Ap- proach	Feasibility Assessment	Time Horizon	Capital Required	Key Barrier	What It Solves
IABs	Moderately challenging	Medium- term (10-20 years)	\$200M-\$2B	Bootstrap funding	Adoption of welfare- improving policies
Wishoc- racy (post- IAB)	Moderate (conditional on IABs)*	Medium- term (5-10 years)	\$50M-\$200M	Requires treaty first	Post-adoption allocation
Futarchy (predic- tion markets)	Very challenging	Long-term (15-25 years)	\$100M-\$500M	Manipula- tion, adoption	Policy → outcome mapping

Ap- proach	Feasibility Assessment	Time Horizon	Capital Required	Key Barrier	What It Solves
Quadratic Vot- ing/Fund- ing	Very challenging	Long-term (10-20 years)	\$50M-\$200M	Constitu- tional barriers	Preference intensity
Algorith- mic Govern- ment	Extremely challenging	Very long-term (30+ years)	\$10B+	Legitimacy, capture risk	Removes human discretion
Charter Cities	Very challenging	Very long-term (20-40 years)	\$1B+	Sovereignty, scale	Competitive governance
NSV Ranking Alone	Infeasible	N/A	\$10M-\$50M	No incentive linkage	Information (ignored)

*Wishocracy becomes feasible only after IABs create political conditions for resource reallocation; otherwise faces same adoption barriers as IABs.

Key insights:

1. **Algorithmic government fails** because the capture problem recurses: whoever controls the algorithm specification controls outcomes. Defense contractors wouldn't disappear; they'd pivot to capturing the algorithm design process.
2. **Futarchy (prediction markets for policy)** is intellectually elegant but faces adoption barriers. Who decides the outcome metrics? How do you prevent market manipulation? IABs could bootstrap futarchy by making politicians willing to adopt prediction-market-based governance.
3. **Charter cities / competitive governance** avoid the adoption problem (exit rather than voice) but face scale and sovereignty constraints.
4. **NSV rankings alone** are infeasible because they lack incentive linkage. The empirical fact motivating this entire paper.

A sequenced implementation strategy works better: 1. **Near-term (0-10 years):** IABs for high-value domains (health, climate, pandemic preparedness) 2. **Medium-term (10-20 years):** Use IAB political capital to pass Wishocracy and futarchy pilots 3. **Long-term (20+ years):** Successful pilots create demand for broader adoption

7.4 What IABs Do Not Solve

IABs solve the *adoption* problem (the binding constraint that prevents welfare-improving policies from being considered) but not the complete resource allocation problem. The four-layer stack addresses this systematically:

Problems IABs alone do not solve:

- **What counts as “public good”?** → **Solved by Layer 0 (Wishocracy for Domain Ranking):** Expert organizations provide BCR data; citizens aggregate via pairwise comparisons to produce democratic domain rankings. This combines expert cost-effectiveness analysis with democratic legitimacy while preventing capture of the scoring methodology.
- **How to allocate within approved domains?** → **Solved by Layer 2:** Wishocracy democratizes allocation via aggregated pairwise preferences, preventing committee capture
- **Which specific projects to fund?** → **Solved by Layer 3:** Domain-specific marketplaces (a decentralized framework for drug assessment for trials, prize markets for breakthroughs) allocate to specific projects
- **How to prevent metric corruption?** → **Partially solved by Layers 0 & 2:** Multiple competing scoring orgs (Layer 0) + democratic allocation (Layer 2) reduce single-point-of-failure risk
- **Democratic legitimacy of wealthy funder influence** → **Solved by Layer 2:** Capital asymmetry (Section 2.3) + Wishocracy (Section 6.2.3) mean millions of retail investors democratize funding, and allocation decisions are made by all citizens via pairwise comparisons, not funders

Problems that remain genuinely unsolved:

1. **Crowding out intrinsic political motivation:** IABs may make politicians more transactional (empirical question requiring evaluation)
2. **Constitutional/legal barriers:** Some jurisdictions may prohibit IAB mechanisms via campaign finance or securities regulation
3. **Transition risk:** Bootstrap phase (12-24 months before proof-of-concept) vulnerable to regulatory counter-attack
4. **Value pluralism vs. utilitarian efficiency:** Even with Wishocracy at Layer 0 allowing citizens to weight domains by their values, some citizens may reject the entire framework of cost-effectiveness analysis in favor of deontological or rights-based approaches

What IABs do solve: The adoption problem. They make it individually rational for politicians to support welfare-improving policies despite concentrated opposition. This is the *binding constraint* in the governance stack. Without solving adoption, better allocation mechanisms (Layers 2-3) cannot be implemented. IABs are needed but not enough for optimal resource allocation.

8 Generalization: IABs for Global Coordination Problems

The IAB architecture is not specific to health policy. It applies to any global coordination problem satisfying three conditions:

1. **Measurable outcomes:** There exists a metric politicians can be scored on
2. **Political control:** Politicians’ actions (votes, treaties, budget allocations) affect outcomes
3. **Diffuse benefits, concentrated costs:** The Olsonian collective action failure applies

These conditions identify domains where the mechanism design approach is applicable. The formal requirements (A1–A5) must be satisfied, with domain-specific adaptations to the scoring function f and payoff calibration.

8.1 The General Template

Any global public good satisfying the above conditions can be “IAB-ified” using the structure in Table 13.

Table 13: General IAB Template for Global Public Goods

Component	Function	Legal Form
Metric	Measures politician-controlled outcomes	Defined by 501(c)(3) research org
Score	Translates actions into public number	Published by independent body
Electoral layer	Rewards high-scorers with campaign support	501(c)(4), PAC, Super PAC
Post-office layer	Reserves prestige positions for high-scorers	Foundations, think tanks

8.2 Candidate Domains

We briefly survey three domains where IABs are applicable; Appendix B provides detailed specifications.

Climate change. Verified emissions reductions (UNFCCC reporting) provide a measurable metric. Politicians can be scored on climate legislation votes, treaty ratification, and budget allocations. The concentrated-diffuse asymmetry is stark: fossil fuel producers face concentrated losses while climate benefits are globally diffuse.

Nuclear disarmament. Verified warhead reductions under START-type treaties provide objective metrics. The scoring function would weight votes on arms control treaties, defense authorization amendments, and non-proliferation funding. Defense contractor opposition creates concentrated costs; security benefits are diffuse.

Pandemic preparedness. WHO Joint External Evaluation scores and pandemic preparedness funding levels are measurable. Health security appropriations votes provide a scoring basis. Pharmaceutical incumbent opposition and budget competition create concentrated resistance; pandemic prevention benefits the global population.

Each domain requires calibration of the opposition cost c_i and the IAB funding levels needed to satisfy Equation 16. The structural requirements remain constant across applications.

8.3 Welfare Accounting of the IAB Mechanism

The paper demonstrates that IABs redirect resources from low-NSV to high-NSV programs, but economists will correctly ask: **what is the net welfare impact accounting for the mechanism’s own costs?**

Mechanism costs (annual, steady-state):

Component	Estimated Annual Cost	Purpose
Scoring organizations	\$10M-\$50M	Research, data, scoring methodology
Independent expenditure campaigns	\$100M-\$500M	Electoral support for high-scorers
Post-office foundations	\$50M-\$200M	Fellowships, advisory positions
Total mechanism overhead	\$160M-\$750M/year	Ongoing operation

Benefits (illustrative, 1% military reallocation treaty):

Benefit Category	Estimated Annual Value	Source
Peace dividend (reduced conflict costs)	\$113.55B	1% × \$11.355T global war costs
Pragmatic clinical trials acceleration	\$24B reallocated × 4.75 BCR	\$114B
Total annual benefits	\$227.55B	Conservative estimate

Net welfare calculation:

$$NSV_{\text{mechanism}} = \text{Benefits} - \text{Costs} = \$227.55\text{B} - \$0.75\text{B} = \$226.8\text{B/year}$$

Benefit-cost ratio of the IAB mechanism itself:

$$BCR_{\text{IAB}} = \frac{\$227.55\text{B}}{\$0.75\text{B}} \approx 303 : 1$$

Even using the high-end cost estimate (\$750M) and conservative benefits (ignoring climate, pandemic preparedness, other applications), **the mechanism's BCR exceeds 300:1**. The overhead cost is 0.3% of benefits, comparable to the expense ratio of efficient index funds (0.03-0.3%).

Sensitivity to failure risk:

If the mechanism has only a 10% chance of achieving treaty passage, expected BCR falls to 30:1, still dramatically positive. If success probability is 50%, expected BCR is 150:1. The mechanism justifies its costs across a wide range of success probabilities.

Comparison to alternative governance reforms:

Reform Mechanism	Estimated Annual Cost	Estimated Benefits	BCR
IABs (this paper)	\$160M-\$750M	\$227B+	303:1 to 1,400:1
Layer 0 only (Wishocracy domain ranking without Layer 1)	\$10M-\$50M	~\$0 (ignored by politicians)	0:1
Futarchy infrastructure	\$500M-\$1B	Uncertain	Un-known
Charter cities	\$1B+	Limited scale	1:1 to 5:1

Key insight: IABs’ welfare accounting is favorable precisely because they solve the *adoption* problem. Better domain prioritization (Layer 0 Wishocracy) costs little but achieves nothing if politicians ignore the rankings. Better allocation mechanisms (futarchy, Layer 2 Wishocracy) are valuable only if adopted, which requires Layer 1 (IABs) first. The mechanism’s overhead is justified by making all other governance improvements politically feasible.

Risks to welfare accounting:

1. **Crowding out intrinsic motivation:** If IABs make politicians more transactional, could reduce welfare from unscored policies (estimated cost: 5-15% of mechanism benefits)
2. **Capture of scoring process:** Layer 0 (Wishocracy for Domain Ranking) mitigates this by requiring capture of ALL expert organizations AND manipulation of millions of citizen preferences. Residual risk if citizen participation is low or manipulable via misinformation (estimated probability: 5-15% over 20 years, down from 10-25% with expert-only rankings)
3. **Regulatory backlash costs:** Legal battles, reputation damage if mechanism is perceived as illegitimate (estimated cost: \$50M-\$200M in defensive spending)

Even accounting for these risks (expected cost: ~\$75M/year, down from ~\$100M/year due to reduced capture risk), the net welfare impact remains strongly positive (BCR > 200:1).

9 Conclusion

Incentive Alignment Bonds represent a new application of mechanism design to democratic governance. By making support for public-good policies incentive-compatible for utility-maximizing politicians, IABs address the collective action failure identified by Olson: concentrated interests systematically defeating diffuse interests in political competition.

In plain English: The mechanism makes it personally profitable for politicians to support welfare-improving policies such as pragmatic clinical trials, rather than harmful or low-value programs.

The contribution is fivefold:

1. **Theoretical:** We formalize political incentive alignment as a mechanism design problem, provide explicit functional forms for politician utility components, and prove conditions for incentive compatibility.
2. **Empirical:** We provide numerical calibration demonstrating that the incentive compatibility condition is achievable with realistic parameters in selected domains (pandemic prepared-

ness, health research) while showing mechanism failure in high-opposition domains (defense reallocation) absent enough funding.

3. **Instrumental:** We define IABs through three primitive properties: investor alignment, politician alignment, and funding from lower-value sources. Together, these create a self-sustaining mechanism for public good production that restores alignment between politician incentives and general welfare.
4. **Practical:** We specify a three-layer architecture (scoring, electoral, post-office) that achieves alignment without violating anti-bribery law, and demonstrate that the bootstrap problem is solvable due to ROI economics and capital asymmetry.
5. **Welfare-Economic:** We provide a comparative criterion ($NSV_{source} < NSV_{target}$) for identifying appropriate funding sources, implying a natural ranking of expenditures by marginal social value. We demonstrate why *specified reallocation*, rather than budget addition, is needed to achieve real resource shifts under soft budget constraints. Welfare accounting shows the mechanism itself has a BCR exceeding 300:1, justifying overhead costs across a wide range of success probabilities.

Important limitations remain. Gaming and metric corruption require ongoing institutional vigilance. The plutocracy objection, that wealthy funders determine priorities, is addressed structurally by Wishocracy (Layer 2), which democratizes allocation decisions via aggregated citizen preferences and prevents advertising-based capture through random pairwise sampling. The capital asymmetry (\$454T household wealth vs. \$5T concentrated interests) combined with retail-accessible securities ensures that funding itself becomes democratized. Unintended consequences on political culture (e.g., crowding out intrinsic motivation) remain empirical questions requiring pilot evaluation. These considerations counsel for careful pilot implementation with rigorous evaluation before scaling.

9.1 The Bootstrap Problem and Regulatory Resistance

The biggest challenge is political: **concentrated interests threatened by IABs will attempt to regulate them out of existence before they can demonstrate effectiveness.**

If IABs successfully redirect even 1% of military spending (\$24B globally), defense contractors will face real losses. They will deploy their existing lobbying infrastructure to ban, restrict, or capture the mechanism. Potential regulatory attacks include:

- **Campaign finance regulation:** Classify IAB-funded independent expenditures as illegal coordination
- **Securities regulation:** Prohibit retail investment in “political outcome bonds”
- **Tax law changes:** Eliminate tax-exempt status for scoring organizations
- **Capture attempts:** Lobby to control the scoring methodology or install friendly board members

The bootstrap paradox: IABs need enough scale to create political incentives strong enough to resist regulatory capture, but concentrated interests will attempt to kill the mechanism before it reaches that scale. It’s the classic Catch-22: you need to be big enough to win, but your enemies will try to kill you while you’re small.

Potential solutions:

1. **First-mover advantage in permissive jurisdictions:** Pilot in countries with strong free speech protections and established independent expenditure precedent (U.S., UK)
2. **Rapid scaling:** Achieve critical mass (\$200M-\$2B deployed) within 18-24 months, before regulatory counter-mobilization
3. **Constitutional protection:** Establish that IAB mechanisms fall under protected political speech (U.S. First Amendment, European Convention on Human Rights Article 10)
4. **Diverse funding sources:** Avoid dependence on any single capital source that could be regulated
5. **International treaty protection:** Once a treaty is adopted, signatories have incentive to protect the mechanism that secured their score-improvement

The mechanism's survival likely depends on achieving proof-of-concept success (one treaty ratified) before incumbent industries can coordinate effective regulatory opposition. This creates a premium on execution speed and strategic sequencing of target domains.

Why the Bootstrap Problem is Solvable: Capital Asymmetry and ROI

The bootstrap problem, raising \$200M-\$2B before concentrated interests mobilize counter-lobbying, appears daunting until we consider two factors: **(1) the massive ROI potential** and **(2) the capital asymmetry** favoring diffuse beneficiaries.

ROI makes bootstrap funding rational for risk-tolerant capital:

As shown in Section 2.2, a \$100M campaign with 20% success probability passing a \$24B/year treaty (10% to investors) yields:

$$E[V] = 0.20 \times \$2.4\text{B}/\text{year} = \$480\text{M}/\text{year}$$

First-year expected ROI: 272%, with the revenue stream continuing in perpetuity. Even with high risk and long timelines, the asymmetric upside makes initial funding economically rational, not philanthropic. This changes the bootstrap calculus: **early investors are not donors making grants; they are rational actors making high-risk, high-return investments.**

Capital asymmetry ensures scalability beyond bootstrap:

The bootstrap phase requires \$200M-\$2B, which is small relative to the aggregate capital available to diffuse beneficiaries:

- **Concentrated opposition** (defense contractors, fossil fuel companies): \$5 trillion market cap, spending \$100M-\$1B/year on lobbying
- **Diffuse beneficiaries** (everyone who benefits from cures, climate stability): \$454 trillion household wealth (Section 2.3)

The 90:1 capital advantage means that **even if concentrated interests attempt counter-lobbying, they face a resource constraint.** If defense contractors allocate \$500M to kill IABs, diffuse beneficiaries can deploy \$5B in response, and still represent only 0.001% of available household wealth. The political change ROI (100-10,000x) that defense contractors exploit becomes accessible to millions of retail investors.

Empirical precedent: Cryptocurrencies raised \$30B+ in ICOs (2017-2018) with far weaker value propositions than 272% expected ROI backed by lobbying economics. Green bonds reached \$500B

outstanding by 2023. If IABs can be structured as retail-accessible securities, the capital mobilization problem is not “can we raise \$200M?” but “can we structure the offering legally?”

The race condition: Concentrated interests can mobilize regulatory opposition within 12-24 months of IAB visibility. But if: 1. Initial \$200M-\$500M is raised from high-risk capital (impact funds, crypto whales, patient billionaires) 2. Rapid deployment achieves proof-of-concept (one scoring cycle, measurable electoral impact) 3. Returns materialize (even partial treaty passage increases bond value) 4. Retail offering opens to millions of diffuse beneficiaries

...then the mechanism reaches escape velocity before regulatory capture becomes feasible. Defense contractors cannot outspend \$50B+ in mobilized retail capital without bankrupting themselves.

Conclusion: The bootstrap problem is solvable because **(a)** ROI economics make initial funding rational for risk-tolerant capital, and **(b)** capital asymmetry ensures that once proof-of-concept succeeds, scaling capital exceeds any plausible counter-lobbying budget. The binding constraint is execution speed, not capital availability.

Future research should address several open questions: What scoring mechanisms are most robust to gaming? How do IABs interact with existing campaign finance institutions? What governance structures best prevent capture of the scoring layer? Can IABs be adapted to non-democratic political systems? What legal strategies best protect the mechanism from regulatory attack? Empirical testing, beginning with single-issue pilot implementations, will be essential to validate the theoretical framework presented here.

If IABs can change political behavior at scale, if they can make supporting measurable public goods the career-maximizing choice for politicians, the architecture becomes available for climate, nuclear risk, pandemic preparedness, and every domain where humanity’s long-term welfare depends on overcoming collective action failures. The mechanism does not require politicians to become better people. It requires only that institutions be designed so rational self-interest points at better outcomes.

References

- AARP. 2023. “Unpaid Caregiver Hours and Economic Value.” *AARP 2023*. <https://www.aarp.org/caregiving/financial-legal/info-2023/unpaid-caregivers-provide-billions-in-care.html>.
- ACLED. 2024. “Active Combat Deaths Annually.” *ACLED: Global Conflict Surged 2024*. <https://acleddata.com/2024/12/12/data-shows-global-conflict-surged-in-2024-the-washington-post/>.
- al., Bentley et. 2023. “NIH Spending on Clinical Trials: 3.3%.” *Bentley Et Al*. <https://www.fiercebiotech.com/biotech/nih-spending-clinical-trials-reached-81b-over-decade>.
- “Annual NIH Economic Impact Report: 2025 Update.” 2025. United for Medical Research. <https://www.unitedformedicalresearch.org/statements/umr-releases-annual-nih-economic-impact-report-2025-update/>.
- Bank, World. “Swiss Military Budget as Percentage of GDP.” *World Bank: Military Expenditure*. <https://data.worldbank.org/indicator/MS.MIL.XPND.GD.ZS?locations=CH>.
- . “Switzerland Vs. US GDP Per Capita Comparison.” *World Bank: Switzerland GDP Per Capita*. <https://data.worldbank.org/indicator/NY.GDP.PCAP.CD?locations=CH>.
- . “World Bank Trade Disruption Cost from Conflict.” *World Bank*. <https://www.worldbank.org/en/topic/trade/publication/trading-away-from-conflict>.
- . 2022. “Infrastructure Investment Economic Multiplier (1.6).” *World Bank: Infrastructure Investment as Stimulus*. <https://blogs.worldbank.org/en/ppps/effectiveness-infrastructure>.

- investment-fiscal-stimulus-what-weve-learned.
- Barro, Robert J., and Charles J. Redlick. 2011. “Macroeconomic Effects from Government Purchases and Taxes.” *The Quarterly Journal of Economics* 126 (1): 51–102. <https://doi.org/10.1093/qje/qjq002>.
- Barta, Zsófia. 2024. “The Logic of Credit: How and Why Credit Rating Agencies Exercise Power Within Global Governance.” *MPIfG Discussion Paper*.
- Based on WHO Global Health Estimates showing 55 million annual deaths / 365 days = 150, 000 per day | WHO. 2024. “150,000 Deaths Per Day from All Causes.” *Based on WHO Global Health Estimates Showing 55 Million Annual Deaths / 365 Days = 150*. <https://www.who.int/news-room/fact-sheets/detail/the-top-10-causes-of-death>.
- Batini, Nicoletta, Mario Di Serio, Matteo Fragetta, Giovanni Melina, and Anthony Waldron. 2021. “Building Back Better: How Big Are Green Spending Multipliers?” *Ecological Economics* 193: 107305. <https://doi.org/10.1016/j.ecolecon.2021.107305>.
- (BIO), Biotechnology Innovation Organization. 2021. “BIO Clinical Development Success Rates 2011-2020.” *Biotechnology Innovation Organization (BIO)*. https://go.bio.org/rs/490-EHZ-999/images/ClinicalDevelopmentSuccessRates2011_2020.pdf.
- BLS. 2024. “Average US Hourly Wage.” *BLS*. <https://www.bls.gov/news.release/pdf/ocwage.pdf>.
- Buchanan, James M., and Gordon Tullock. 1962. *The Calculus of Consent: Logical Foundations of Constitutional Democracy*. Ann Arbor: University of Michigan Press.
- budgets:, See component country. “Global Government Medical Research Spending (\$67.5B, 2023–2024).” *See Component Country Budgets: NIH Budget*. [#nih-budget-fy2025](https://www.nih.gov/budget).
- Bureau, US Census. “Historical World Population Estimates.” *US Census Bureau*. <https://www.census.gov/data/tables/time-series/demo/international-programs/historical-est-worldpop.html>.
- . 2024. “Number of Registered or Eligible Voters in the u.s.” *US Census Bureau*. <https://www.census.gov/newsroom/press-releases/2025/2024-presidential-election-voting-registration-tables.html>.
- CAN, ACS. “Clinical Trial Patient Participation Rate.” *ACS CAN: Barriers to Clinical Trial Enrollment*. <https://www.fightcancer.org/policy-resources/barriers-patient-enrollment-therapeutic-clinical-trials-cancer>.
- Cardiology, Int’l Journal of. 2050. “Annual Global Economic Burden of Heart Disease.” *Int’l Journal of Cardiology: Global Heart Failure Burden02238-9/Abstract*. [https://www.internationaljournalofcardiology.com/article/S0167-5273\(13](https://www.internationaljournalofcardiology.com/article/S0167-5273(13)
- Care, Diabetes. “Annual Global Economic Burden of Diabetes.” *Diabetes Care: Global Economic Burden*. <https://diabetesjournals.org/care/article/41/5/963/36522/Global-Economic-Burden-of-Diabetes-in-Adults>.
- CDC. “U.s. Chronic Disease Healthcare Spending.” *CDC*. <https://www.cdc.gov/chronic-disease/data-research/facts-stats/index.html>.
- . 2017. “Childhood Vaccination (US) ROI.” *CDC*. <https://www.cdc.gov/mmwr/preview/mmwrhtml/mm6316a4.htm>.
- C&EN. 2025. “Annual Number of New Drugs Approved Globally: 50.” *C&EN*. <https://cen.acs.org/pharmaceuticals/50-new-drugs-received-FDA/103/i2>.
- CGDev. 2024. “UNHCR Average Refugee Support Cost.” *CGDev*. <https://www.cgdev.org/blog/costs-hosting-refugees-oecd-countries-and-why-uk-outlier>.
- CNBC. 2025. “Warren Buffett’s Career Average Investment Return.” *CNBC*. <https://www.cnbc.com/2025/05/05/warren-buffetts-return-tally-after-60-years-5502284percent.html>.
- Code, United States. “18 u.s. Code 201 - Bribery of Public Officials and Witnesses.” <https://www.law.cornell.edu/uscode/text/18/201>.
- Committee, Nobel Prize. 2007. “The Sveriges Riksbank Prize in Economic Sciences in Memory of

Alfred Nobel 2007.” <https://www.nobelprize.org/prizes/economic-sciences/2007/summary/>.

Conservation Voters, League of. 2024. “About the Scorecard.” <https://scorecard.lcv.org/methodology>.

“Credit Rating Agencies and Sovereign Debt: Four Proposals to Support Achievement of the SDGs.” 2023. 131. United Nations Department of Economic; Social Affairs. <https://desapublications.un.org/policy-briefs/un-desa-policy-brief-no-131>.

CSDD, Tufts. “Cost of Drug Development.”

CSIS. “Smallpox Eradication ROI.” *CSIS*. <https://www.csis.org/analysis/smallpox-eradication-model-global-cooperation>.

Data, Our World in. 2024. “Terror Attack Deaths (8,300 Annually).” *Our World in Data: Terrorism*. <https://ourworldindata.org/terrorism>.

Deloitte. 2025. “Pharmaceutical r&d Return on Investment (ROI).” *Deloitte: Measuring Pharmaceutical Innovation 2025*. <https://www.deloitte.com/ch/en/Industries/life-sciences-health-care/research/measuring-return-from-pharmaceutical-innovation.html>.

Discovery, Nature Reviews Drug. 2016. “Drug Trial Success Rate from Phase i to Approval.” *Nature Reviews Drug Discovery: Clinical Success Rates*. <https://www.nature.com/articles/nrd.2016.136>.

DOT. 2024. “DOT Value of Statistical Life (\$13.6M).” *DOT: VSL Guidance 2024*. <https://www.transportation.gov/office-policy/transportation-policy/revised-departmental-guidance-on-valuation-of-a-statistical-life-in-economic-analysis>.

Drugs.com. “FDA Drug Approval Timeline.” *Drugs.com: FDA Drug Approval Process*. <https://www.drugs.com/fda-approval-process.html>.

EPI. “Education Investment Economic Multiplier (2.1).” *EPI: Public Investments Outside Core Infrastructure*. <https://www.epi.org/publication/bp348-public-investments-outside-core-infrastructure/>.

———. 2024. “CEO Compensation.” *EPI*. <https://www.epi.org/blog/ceo-pay-increased-in-2024-and-is-now-281-times-that-of-the-typical-worker-new-epi-landing-page-has-all-the-details/>.

“Estimated Impact of the American Recovery and Reinvestment Act on Employment and Economic Output in 2014.” 2015. Congressional Budget Office. <https://www.cbo.gov/publication/49958>.

estimates, Industry. “Clinical Trial Abandonment.”

FDA. “FDA GRAS List Count (570-700).” *FDA*. <https://www.fda.gov/food/generally-recognized-safe-gras/gras-notice-inventory>.

———. “FDA-Approved Prescription Drug Products (20,000+).” *FDA*. <https://www.fda.gov/media/143704/download>.

“Federal Lobbying Set New Record in 2024.” 2025. OpenSecrets. <https://www.opensecrets.org/news/2025/02/federal-lobbying-set-new-record-in-2024/>.

Fowler, Anthony, Haritz Garro, and Jörg L. Spenkuch. 2020. “Quid Pro Quo? Corporate Returns to Campaign Contributions.” *The Journal of Politics* 82 (3): 844–58.

GAO. “Annual Cost of u.s. Sugar Subsidies.” *GAO: Sugar Program*. <https://www.gao.gov/products/gao-24-106144>.

———. 2025. “95% of Diseases Have No Effective Treatment.” *GAO*. <https://www.gao.gov/products/gao-25-106774>.

Gilens, Martin, and Benjamin I. Page. 2014. “Testing Theories of American Politics: Elites, Interest Groups, and Average Citizens.” *Perspectives on Politics* 12 (3): 564–81. <https://doi.org/10.1017/S1537592714001595>.

GiveWell. “Cost Per DALY for Deworming Programs.” <https://www.givewell.org/international/technical/programs/deworming/cost-effectiveness>.

———. “GiveWell Cost Per Life Saved for Top Charities (2024).” *GiveWell: Top Charities*. <https://www.givewell.org/charities/top-charities>.

- “Global Health Estimates: Life Expectancy and Leading Causes of Death and Disability.” 2024. World Health Organization. <https://www.who.int/data/gho/data/themes/mortality-and-global-health-estimates>.
- “Global Wealth Report 2023.” 2023. UBS; Credit Suisse. <https://www.ubs.com/global/en/family-office-uhnw/reports/global-wealth-report-2023.html>.
- Goodhart, Charles A. E. 1984. “Problems of Monetary Management: The UK Experience.” *Monetary Theory and Practice*, 91–121.
- “Halftime for SDGs: Child Immunization.” 2023. Copenhagen Consensus. <https://copenhagenconsensus.com/publication/halftime-sdgs-child-immunization>.
- Health, Value in. “Average Lifetime Revenue Per Successful Drug.” *Value in Health: Sales Revenues for New Therapeutic Agents* 02754-2/Fulltext). [https://www.valueinhealthjournal.com/article/S1098-3015\(24](https://www.valueinhealthjournal.com/article/S1098-3015(24).
- Heckman, James J., Seong Hyeok Moon, Rodrigo Pinto, Peter A. Savelyev, and Adam Yavitz. 2010. “The Rate of Return to the HighScope Perry Preschool Program.” *Journal of Public Economics* 94 (1-2): 114–28. <https://doi.org/10.1016/j.jpubeco.2009.11.001>.
- Hevenstone, Debra. 2023. “The Impact of Social Impact Bond Financing.” *Public Administration Review* 83 (4): 821–36. <https://doi.org/10.1111/puar.13631>.
- Hurwicz, Leonid. 1972. “On Informationally Decentralized Systems.” *Decision and Organization*, 297–336.
- ICER. 2024. “Value Per QALY (Standard Economic Value).” *ICER*. <https://icer.org/wp-content/uploads/2024/02/Reference-Case-4.3.25.pdf>.
- ICRC. 1997. “International Campaign to Ban Landmines (ICBL) - Ottawa Treaty (1997).” *ICRC*. <https://www.icrc.org/en/doc/resources/documents/article/other/57jpjn.htm>.
- IHME Global Burden of Disease (2.55B DALYs), Calculated from, and global GDP per capita valuation. “\$109 Trillion Annual Global Disease Burden.”
- “IMF Fossil Fuel Subsidies Data: 2023 Update.” 2023. Black, Simon; Parry, Ian; Vernon, Nate. <https://www.imf.org/en/Topics/climate-change/energy-subsidies>.
- Institute, Cato. “Chance of Dying from Terrorism Statistic.” *Cato Institute: Terrorism and Immigration Risk Analysis*. <https://www.cato.org/policy-analysis/terrorism-immigration-risk-analysis>.
- Institute, Integrated Benefits. 2024. “Chronic Illness Workforce Productivity Loss.” *Integrated Benefits Institute 2024*. <https://www.ibiweb.org/resources/chronic-conditions-in-the-us-workforce-prevalence-trends-and-productivity-impacts>.
- Institute, Manhattan. “RECOVERY Trial 82× Cost Reduction.” *Manhattan Institute: Slow Costly Trials*. <https://manhattan.institute/article/slow-costly-clinical-trials-drag-down-biomedical-breakthroughs>.
- Laffont, Jean-Jacques, and Jean Tirole. 1993. *A Theory of Incentives in Procurement and Regulation*. MIT Press.
- “Lessons Learned from the Planning and Early Implementation of the Social Impact Bond at HMP Peterborough.” 2011. Disley, Emma; Rubin, Jennifer; Scraggs, Emily; Burrowes, Nina; Culley, Deirdre.
- Liebman, Jeffrey B. 2011. “Social Impact Bonds: A Guide for State and Local Governments.” *Harvard Kennedy School Social Impact Bond Technical Assistance Lab*. <https://www.hks.harvard.edu/centers/mrcbg/publications/awp>.
- literature, Educational psychology. “Average Reading Speed.”
- Lohmann, Susanne. 1998. “An Information Rationale for the Power of Special Interests.” *American Political Science Review* 92 (4): 809–27. <https://www.cambridge.org/core/journals/american-political-science-review/article/abs/an-information-rationale-for-the-power-of-special->

- interests/2321FE7C64F1799F58A69BB5A1FE594D.
- “March of the Four-Stars: The Role of Retired Generals and Admirals in the Arms Industry.” 2023. Hartung, William D.; Fisher, Dillon. <https://quincyinst.org/research/march-of-the-four-stars-the-role-of-retired-generals-and-admirals-in-the-arms-industry/>.
- Maskin, Eric. 1999. “Nash Equilibrium and Welfare Optimality.” *Review of Economic Studies* 66 (1): 23–38.
- “Measuring the Success of Impact Bonds.” 2017. Gustafsson-Wright, Emily; Gardiner, Sophie; Putcha, Vidya. <https://www.brookings.edu/articles/measuring-the-success-of-impact-bonds/>.
- Medicine, JAMA Internal. “Phase 3 Cost Per Patient.” *JAMA Internal Medicine*. <https://jamanetwork.com/journals/jamainternalmedicine/fullarticle/2702287>.
- Medicine, Nature. 2024. “Drug Repurposing Rate (30%).” *Nature Medicine*. <https://www.nature.com/articles/s41591-024-03233-x>.
- Mercatus. “Military Spending Economic Multiplier (0.6).” *Mercatus: Defense Spending and Economy*. <https://www.mercatus.org/research/research-papers/defense-spending-and-economy>.
- Milgrom, Paul. 2004. *Putting Auction Theory to Work*. Cambridge University Press.
- MMWR, CDC. 1994. “Childhood Vaccination Economic Benefits.” *CDC MMWR*. <https://www.cdc.gov/mmwr/volumes/73/wr/mm7331a2.htm>.
- “More Money, Less Transparency: A Decade Under Citizens United.” 2020. OpenSecrets. <https://www.opensecrets.org/news/reports/a-decade-under-citizens-united>.
- Myerson, Roger B. 1979. “Incentive Compatibility and the Bargaining Problem.” *Econometrica* 47 (1): 61–73.
- NCBI, FDA Study via. “Trial Costs, FDA Study.” *FDA Study via NCBI*. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6248200/>.
- News, UN. 2014. “Clean Water & Sanitation (LMICs) ROI.” *UN News*. <https://news.un.org/en/story/2014/11/484032>.
- NHGRI. 2003. “Human Genome Project and CRISPR Discovery.” *NHGRI*. <https://www.genome.gov/11006929/2003-release-international-consortium-completes-hgp>.
- NIH. 2015. “Antidepressant Clinical Trial Exclusion Rates.” *Zimmerman Et Al*. <https://pubmed.ncbi.nlm.nih.gov/26276679/>.
- Numbers, Think by. “Post-1962 Drop in New Drug Approvals.” *Think by Numbers: How Many Lives Does FDA Save?* <https://thinkbynumbers.org/health/how-many-net-lives-does-the-fda-save/>.
- . 1962. “Pre-1962 Drug Development Costs and Timeline.” *Think by Numbers: How Many Lives Does FDA Save?* <https://thinkbynumbers.org/health/how-many-net-lives-does-the-fda-save/>.
- . 1966. “Pre-1962 Physician-Led Clinical Trials.” *Think by Numbers: How Many Lives Does FDA Save?* <https://thinkbynumbers.org/health/how-many-net-lives-does-the-fda-save/>.
- . 2021. “Lost Human Capital Due to War (\$270B Annually).” <https://thinkbynumbers.org/military/war/the-economic-case-for-peace-a-comprehensive-financial-analysis/%3E>.
- Olson, Mancur. 1965. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Cambridge, MA: Harvard University Press.
- Oncology, JAMA. 2020. “Annual Global Economic Burden of Cancer.” *JAMA Oncology: Global Cost 2020-2050*. <https://jamanetwork.com/journals/jamaoncology/fullarticle/2801798>.
- ONE, PLOS. 2010. “Cost Per DALY for Vitamin a Supplementation.” *PLOS ONE: Cost-Effectiveness of "Golden Mustard" for Treating Vitamin A Deficiency in India (2010)*. <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0012046>.
- One, PLOS. 2019. “Health and Quality of Life of Thalidomide Survivors as They Age.” *PLOS One*. <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0210222>.
- OpenSecrets. “Lobbyist Statistics for Washington d.c.” *OpenSecrets: Lobbying in US*. <https://www.opensecrets.org/lobby>.

- [//en.wikipedia.org/wiki/Lobbying_in_the_United_States](https://en.wikipedia.org/wiki/Lobbying_in_the_United_States).
- . 2024. “Lobbying Spend (Defense).” *OpenSecrets*. <https://www.opensecrets.org/federal-lobbying/industries/summary?cycle=2024/&id=D>.
- Oren Cass, Manhattan Institute. 2023. “RECOVERY Trial Cost Per Patient.” *Oren Cass*. <https://manhattan.institute/article/slow-costly-clinical-trials-drag-down-biomedical-breakthroughs>.
- Organization, World Health. 2022. “Mental Health Global Burden.” *World Health Organization*. <https://www.who.int/news/item/28-09-2001-the-world-health-report-2001-mental-disorders-affect-one-in-four-people>.
- . 2024. “WHO Global Health Estimates 2024.” *World Health Organization*. <https://www.who.int/data/gho/data/themes/mortality-and-global-health-estimates>.
- PMC. PMC. <https://pmc.ncbi.nlm.nih.gov/articles/PMC5193154/>.
- . 2012. “Contribution of Smoking Reduction to Life Expectancy Gains.” *PMC: Benefits Smoking Cessation Longevity*. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1447499/>.
- . 2022. “Healthcare Investment Economic Multiplier (1.8).” *PMC: California Universal Health Care*. <https://pmc.ncbi.nlm.nih.gov/articles/PMC5954824/>.
- . 2023. “Only 12% of Human Interactome Targeted.” *PMC*. <https://pmc.ncbi.nlm.nih.gov/articles/PMC10749231/>.
- PubMed. “Psychological Impact of War Cost (\$100B Annually).” *PubMed: Economic Burden of PTSD*. <https://pubmed.ncbi.nlm.nih.gov/35485933/>.
- Reardon, Sara. 2014. “Lobbying Sways NIH Grants.” *Nature* 515: 19. <https://www.nature.com/articles/515019a>.
- Report, IQVIA. “Global Trial Capacity.” *IQVIA Report: Clinical Trial Subjects Number Drops Due to Decline in COVID-19 Enrollment*. <https://gmdpacademy.org/news/iqvia-report-clinical-trial-subjects-number-drops-due-to-decline-in-covid-19-enrollment/>.
- Research, ProRelix. “Clinical Trial Cost Per Patient (Traditional Phase III).” *ProRelix Research*. <https://prorelixresearch.com/phase-by-phase-clinical-trial-costs-what-every-sponsor-needs-to-know/>.
- Research, and Markets. 2024. “Global Clinical Trials Market 2024.” *Research and Markets*. <https://www.globenewswire.com/news-release/2024/04/19/2866012/0/en/Global-Clinical-Trials-Market-Research-Report-2024-An-83-16-Billion-Market-by-2030-AI-Machine-Learning-and-Blockchain-will-Transform-the-Clinical-Trials-Landscape.html>.
- Reynolds, Arthur J., Judy A. Temple, Barry A. White, Suh-Ruu Ou, and Dylan L. Robertson. 2011. “Age 26 Cost-Benefit Analysis of the Child-Parent Center Early Education Program.” *Child Development* 82 (1): 379–404. <https://doi.org/10.1111/j.1467-8624.2010.01563.x>.
- Roth, Alvin E. 2002. “The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics.” *Econometrica* 70 (4): 1341–78.
- School, Harvard Kennedy. 2020. “3.5% Participation Tipping Point.” *Harvard Kennedy School*. <https://www.hks.harvard.edu/centers/carr/publications/35-rule-how-small-minority-can-change-world>.
- ScienceDaily. 2015. “Global Prevalence of Chronic Disease.” *ScienceDaily: GBD 2015 Study*. <https://www.sciencedaily.com/releases/2015/06/150608081753.htm>.
- . 2020. “Medical Research Lives Saved Annually (4.2 Million).” *ScienceDaily: Physical Activity Prevents 4M Deaths*. <https://www.sciencedaily.com/releases/2020/06/200617194510.htm>.
- Service, Internal Revenue. “Exemption Requirements - 501(c)(3) Organizations.” <https://www.irs.gov/charities-non-profits/charitable-organizations/exemption-requirements-501c3-organizations>.
- Shepherd, Michael E., and Hye Young You. 2019. “Exit Strategy: Career Concerns and Revolving Doors in Congress.” *American Political Science Review* 114 (1): 270–84. <https://doi.org/10.1017/XPS.2019.1>.

- 1017/S000305541900067X.
- SIPRI. 2016. “36:1 Disparity Ratio of Spending on Weapons over Cures.” *SIPRI: Military Spending*. <https://www.sipri.org/commentary/blog/2016/opportunity-cost-world-military-spending>.
- . 2025. “Global Military Spending (\$2.72T, 2024).” *SIPRI*. <https://www.sipri.org/publications/2025/sipri-fact-sheets/trends-world-military-expenditure-2024>.
- Smithberger, Mandy, and William Hartung. 2021. “Top Defense Firms See 2Treturnon1B Investment in Afghan War.” *Responsible Statecraft*. <https://responsiblestatecraft.org/2021/09/02/top-defense-firms-see-2t-return-on-1b-investment-in-afghan-war/>.
- SofproMed. “Phase 3 Cost Per Trial Range.” *SofproMed*. <https://www.sofpromed.com/how-much-does-a-clinical-trial-cost>.
- Sood, Neeraj, Rachel Ribero, Megan Ryan, and Karen Van Nuys. 2023. “The Revolving Door in Health Care Regulation.” *Health Affairs* 42 (9): 1284–91. <https://doi.org/10.1377/hlthaff.2023.00418>.
- Spenkuch, Jörg L., and David Toniatti. 2018. “Does Campaign Spending Affect Election Outcomes? New Evidence from Transaction-Level Disbursement Data.” *The Journal of Politics* 82 (4): 1502–15. <https://doi.org/10.1086/708646>.
- Statista. 2024. “US Military Budget as Percentage of GDP.” *Statista*. <https://www.statista.com/statistics/262742/countries-with-the-highest-military-spending/>.
- Stratmann, Thomas. 2002. “Can Special Interests Buy Congressional Votes? Evidence from Financial Services Legislation.” *The Journal of Law and Economics* 45 (2): 345–73. <https://doi.org/10.1086/340091>.
- Suisse/UBS, Credit. 2023. “Credit Suisse Global Wealth Report 2023.” *Credit Suisse/UBS*. <https://www.ubs.com/global/en/family-office-uhnw/reports/global-wealth-report-2023.html>.
- “Trends in World Military Expenditure, 2023.” 2024. Stockholm International Peace Research Institute. <https://www.sipri.org/publications/2024/sipri-fact-sheets/trends-world-military-expenditure-2023>.
- Trials. “Patient Willingness to Participate in Clinical Trials.” *Trials: Patients’ Willingness Survey*. <https://trialsjournal.biomedcentral.com/articles/10.1186/s13063-015-1105-3>.
- Trials, Applied Clinical. *Applied Clinical Trials*. <https://www.appliedclinicaltrialsonline.com/view/sizing-clinical-research-market>.
- UCDP. “State Violence Deaths Annually.” *UCDP: Uppsala Conflict Data Program*. <https://ucdp.uu.se/>.
- UN. 2022. “Global Population Reaches 8 Billion.” *UN: World Population 8 Billion Nov 15 2022*. <https://www.un.org/en/desa/world-population-reach-8-billion-15-november-2022>.
- United States, Supreme Court of the. 1999. “United States v. Sun-Diamond Growers of California.”
- . 2010. “Citizens United v. Federal Election Commission.”
- . 2014. “McCutcheon v. Federal Election Commission.”
- . 2016. “McDonnell v. United States.”
- VA. 2026. “Veteran Healthcare Cost Projections.” *VA*. <https://department.va.gov/wp-content/uploads/2025/06/2026-Budget-in-Brief.pdf>.
- Vaccines, MDPI. 2024. “Measles Vaccination ROI.” *MDPI Vaccines*. <https://www.mdpi.com/2076-393X/12/11/1210>.
- via, Direct analysis. “ClinicalTrials.gov Cumulative Enrollment Data (2025).” *Direct Analysis via ClinicalTrials.gov API V2*. <https://clinicaltrials.gov/data-api/api>.
- War, Brown Watson Costs of. “Environmental Cost of War (\$100B Annually).” *Brown Watson Costs of War: Environmental Cost*. <https://watson.brown.edu/costsofwar/costs/social/environment>.
- WHO. 2019a. “Annual Global Economic Burden of Alzheimer’s and Other Dementias.” *WHO: Dementia Fact Sheet*. <https://www.who.int/news-room/fact-sheets/detail/dementia>.

- . 2019b. “ICD-10 Code Count (14,000).” *WHO*. <https://icd.who.int/browse10/2019/en>.
- . 2019c. “Polio Vaccination ROI.” *WHO*. <https://www.who.int/news-room/feature-stories/detail/sustaining-polio-investments-offers-a-high-return>.
- Wikipedia. “Longevity Escape Velocity (LEV) - Maximum Human Life Extension Potential.” *Wikipedia: Longevity Escape Velocity*. https://en.wikipedia.org/wiki/Longevity_escape_velocity.
- . “Thalidomide Scandal: Worldwide Cases and Mortality.” *Wikipedia*. https://en.wikipedia.org/wiki/Thalidomide_scandal.
- . 2020. “US Military Spending Reduction After WWII.” *Wikipedia*. https://en.wikipedia.org/wiki/Demobilization_of_United_States_Armed_Forces_after_World_War_II.
- Wishonia. 2024. “Wishocracy: A Decentralized Semi-Autonomous Todo List for Humanity.” <https://github.com/wishonia/wishonia>.

Appendix A: Formal Proofs

This appendix provides complete proofs of the propositions stated in Section 3.

9.2 Proof of Proposition 1 (Sufficient Condition for Incentive Compatibility)

Statement. Under assumptions A1–A5, if the score gain from supporting policy class \mathcal{P} is $\Delta\theta > 0$, and

$$\alpha_i \cdot \Delta P_i + \beta_i \cdot \Delta Y_i + \gamma_i \cdot \Delta S_i > c_i$$

then $a_i = 1$ is the unique best response for politician i .

Proof.

Step 1: Action space. By A5, politician i faces a binary choice $a_i \in \{0, 1\}$.

Step 2: Information structure. By A2, voting records are publicly observable and the scoring function $f : \text{VoteRecord} \rightarrow \mathbb{R}$ is common knowledge. Hence, politician i can compute:

$$\begin{aligned}\theta_i(a_i = 1) &= \theta_i^0 + \Delta\theta \\ \theta_i(a_i = 0) &= \theta_i^0\end{aligned}$$

where θ_i^0 is the current score and $\Delta\theta > 0$ is the score increment from supporting \mathcal{P} .

Step 3: Payoff functions. By A3, the IAB mechanism credibly commits to score-dependent payoffs. The politician can therefore compute the payoff differentials:

For the electoral component, by Equation 13:

$$\Delta P_i = P_i(\theta_i^0 + \Delta\theta) - P_i(\theta_i^0) = \delta \cdot \Delta\theta + \epsilon_i \cdot [I_i(\theta_i^0 + \Delta\theta) - I_i(\theta_i^0)]$$

For the post-office component, by Equation 15:

$$\Delta Y_i = Y(\tau(\theta_i^0 + \Delta\theta)) - Y(\tau(\theta_i^0))$$

For the legacy component:

$$\Delta S_i = S_i(\theta_i^0 + \Delta\theta) - S_i(\theta_i^0)$$

Step 4: Utility comparison. By A1, politician i maximizes expected utility U_i as defined in Equation 3. The change in utility from choosing $a_i = 1$ versus $a_i = 0$ is:

$$\Delta U_i = U_i(a_i = 1) - U_i(a_i = 0) = \alpha_i \cdot \Delta P_i + \beta_i \cdot \Delta Y_i + \gamma_i \cdot \Delta S_i - c_i$$

where $c_i > 0$ represents the concentrated costs from opposition interests.

Step 5: Optimality. When $\alpha_i \cdot \Delta P_i + \beta_i \cdot \Delta Y_i + \gamma_i \cdot \Delta S_i > c_i$, we have $\Delta U_i > 0$. Since $a_i \in \{0, 1\}$ and choosing $a_i = 1$ yields strictly higher utility, $a_i = 1$ is the unique best response.

Step 6: Uniqueness. The best response is unique because the action space is finite (binary) and $\Delta U_i > 0$ implies strict preference. No indifference exists. \square

9.3 Proof of Corollary 1 (Funding Threshold)

Statement. Under A4, there exists a funding level \bar{F} such that for all $F > \bar{F}$, Equation 16 holds for all politicians with $c_i < \bar{c}$ for some threshold $\bar{c}(F)$ increasing in F .

Proof.

Step 1: Funding dependence. By construction of the IAB mechanism, both ΔP_i and ΔY_i are increasing functions of the total IAB funding F :

- Electoral layer: Higher F implies larger independent expenditure capacity M in Equation 14, increasing ΔP_i
- Post-office layer: Higher F implies more fellowship positions and higher salaries, increasing ΔY_i

Let $\Delta P_i(F)$ and $\Delta Y_i(F)$ denote these functions, with $\frac{\partial \Delta P_i}{\partial F} > 0$ and $\frac{\partial \Delta Y_i}{\partial F} > 0$.

Step 2: Define the benefit function. Let:

$$B_i(F) = \alpha_i \cdot \Delta P_i(F) + \beta_i \cdot \Delta Y_i(F) + \gamma_i \cdot \Delta S_i$$

Since $\alpha_i, \beta_i > 0$ and both $\Delta P_i(F)$ and $\Delta Y_i(F)$ are increasing in F , we have $\frac{\partial B_i}{\partial F} > 0$.

Step 3: Threshold construction. For any $\bar{c} > 0$, define:

$$\bar{F}(\bar{c}) = \inf\{F : B_i(F) \geq \bar{c} \text{ for all } i\}$$

This is well-defined because $B_i(F) \rightarrow \infty$ as $F \rightarrow \infty$ (under A4, the mechanism can scale payoffs arbitrarily with funding).

Step 4: Monotonicity. For $F > \bar{F}(\bar{c})$, we have $B_i(F) > \bar{c}$ for all i . Hence, for any politician with $c_i < \bar{c}$:

$$B_i(F) > \bar{c} > c_i$$

which is precisely the incentive compatibility condition Equation 16.

Step 5: Increasing threshold. The function $\bar{c}(F) = \min_i B_i(F)$ is increasing in F because each $B_i(F)$ is increasing. \square

9.4 Proof of Proposition 2 (Multiple Equilibria Without IABs)

Statement. Under A1–A2 and A5, without the IAB mechanism, the game among N politicians has at least two pure strategy Nash equilibria: (i) the all-defect equilibrium $(0, \dots, 0)$, and (ii) potentially the all-cooperate equilibrium $(1, \dots, 1)$ if coordination is feasible. The all-defect equilibrium is risk-dominant when $c_i > \epsilon$ for all i .

Proof.

Step 1: Payoff structure without IABs. By Equation 11, the payoff gain from choosing $a_i = 1$ versus $a_i = 0$ is:

$$\Delta U_i^{\text{pre-IAB}} = \epsilon - c_i$$

where ϵ represents diffuse voter approval and c_i represents concentrated opposition costs.

Step 2: All-defect equilibrium. Consider the strategy profile $(a_1, \dots, a_N) = (0, \dots, 0)$. For any politician i , unilateral deviation to $a_i = 1$ yields:

$$\Delta U_i = \epsilon - c_i < 0 \quad (\text{by assumption } c_i > \epsilon)$$

Since no politician can improve utility by unilateral deviation, $(0, \dots, 0)$ is a Nash equilibrium.

Step 3: Potential cooperation equilibrium. If politicians can coordinate (through party discipline, repeated game dynamics, or reputation mechanisms), the all-cooperate profile $(1, \dots, 1)$ may be sustainable. In this case:

- Collective benefits may exceed individual costs if policy success generates concentrated rewards (e.g., historical legacy, party brand value)
- Defection can be punished in subsequent interactions

However, this equilibrium requires coordination mechanisms external to the single-shot game.

Step 4: Risk dominance. An equilibrium is risk-dominant if it is the best response to the belief that opponents choose each strategy with equal probability. Under uniform mixing, the expected payoff from $a_i = 0$ exceeds that from $a_i = 1$ when $c_i > \epsilon$, making $(0, \dots, 0)$ risk-dominant. \square

9.5 Proof of Proposition 3 (Equilibrium Selection With IABs)

Statement. Under A1–A5, if the IAB mechanism is funded such that Equation 16 holds for all i , then $(1, \dots, 1)$ is the unique Nash equilibrium.

Proof.

Step 1: Dominant strategy. By Proposition 1, when Equation 16 holds for politician i , choosing $a_i = 1$ is the unique best response regardless of other politicians' choices. This is because the payoff comparison depends only on i 's own score change and the pre-committed IAB payoff rules, not on the actions of other politicians.

Formally, $\Delta U_i > 0$ holds for all strategy profiles $(a_{-i}) \in \{0, 1\}^{N-1}$ of other politicians.

Step 2: Strict dominance. Since $a_i = 1$ yields strictly higher utility than $a_i = 0$ for all a_{-i} , the strategy $a_i = 1$ strictly dominates $a_i = 0$.

Step 3: Unique equilibrium. When every player has a strictly dominant strategy, the profile of dominant strategies is the unique Nash equilibrium. Since $a_i = 1$ is strictly dominant for all $i \in \{1, \dots, N\}$, the profile $(1, \dots, 1)$ is the unique Nash equilibrium.

Step 4: Uniqueness verification. No other strategy profile can be a Nash equilibrium because at any profile containing $a_i = 0$ for some i , politician i has a profitable deviation to $a_i = 1$. \square

Appendix B: Detailed Application Specifications

This appendix provides detailed IAB specifications for the candidate domains discussed in Section 6.

9.6 B.1 Climate Change

Policy objective. Reduce global greenhouse gas emissions in line with Paris Agreement targets (limiting warming to 1.5–2°C).

Scoring metric. The Climate Leadership Score θ_i^C is computed as a weighted average:

$$\theta_i^C = w_1 \cdot V_i^{\text{legislation}} + w_2 \cdot V_i^{\text{treaty}} + w_3 \cdot V_i^{\text{budget}} + w_4 \cdot V_i^{\text{oversight}}$$

where:

- $V_i^{\text{legislation}}$: Voting record on climate legislation (carbon pricing, renewable standards, efficiency mandates)
- V_i^{treaty} : Support for international climate agreements
- V_i^{budget} : Votes on clean energy appropriations and fossil fuel subsidy reform
- $V_i^{\text{oversight}}$: Participation in climate-related oversight activities

Weights w_j sum to 1 and are calibrated to policy impact.

Electoral layer specification.

Score Tier	Independent Expenditure Rule
$\theta_i^C \geq 80$	Full support: +\$M per competitive race
$60 \leq \theta_i^C < 80$	Neutral: no expenditure
$\theta_i^C < 60$	Opposition: −\$M per competitive race

Post-office layer specification.

- Tier 1 ($\theta_i^C \geq 75$): Eligibility for clean energy foundation boards, climate advisory positions, green bank directorships
- Tier 2 ($60 \leq \theta_i^C < 75$): Standard think tank positions
- Tier 3 ($\theta_i^C < 60$): Fossil fuel industry positions remain available but climate-aligned positions closed

Calibration notes. Fossil fuel industry opposition spending averages \$100–500M annually on federal elections. IAB funding at \$1B+ annually would satisfy Equation 16 for most legislators in competitive districts.

9.7 B.2 Nuclear Disarmament

Policy objective. Verified reduction of global nuclear arsenals and strengthened non-proliferation.

Scoring metric. The Nuclear Safety Score θ_i^N :

$$\theta_i^N = w_1 \cdot V_i^{\text{treaty}} + w_2 \cdot V_i^{\text{NDAA}} + w_3 \cdot V_i^{\text{nonprolif}} + w_4 \cdot V_i^{\text{oversight}}$$

where:

- V_i^{treaty} : Votes on arms control treaty ratification (New START, CTBT, etc.)
- V_i^{NDAA} : Votes on National Defense Authorization Act amendments affecting nuclear posture
- $V_i^{\text{nonprolif}}$: Support for non-proliferation funding and diplomacy
- $V_i^{\text{oversight}}$: Engagement with nuclear policy oversight

Electoral layer specification. Peace-focused PACs commit to independent expenditure rules analogous to the climate case, with thresholds calibrated to the nuclear policy domain.

Post-office layer specification.

- Tier 1: Arms control advisory positions, security fellowships at peace-oriented institutions, disarmament diplomacy roles
- Tier 2: General foreign policy positions
- Tier 3: Defense contractor positions remain available but arms control positions closed

Calibration notes. Defense contractor opposition is substantial (over \$100M annually in lobbying). However, the concentrated-diffuse asymmetry is extreme: nuclear war risks affect all humanity. IAB funding levels comparable to current defense lobbying would create countervailing incentives.

9.8 B.3 Pandemic Preparedness

Policy objective. Strengthen global health security infrastructure to prevent and respond to pandemic threats.

Scoring metric. The Pandemic Readiness Score θ_i^P :

$$\theta_i^P = w_1 \cdot V_i^{\text{funding}} + w_2 \cdot V_i^{\text{IHR}} + w_3 \cdot V_i^{\text{surge}} + w_4 \cdot V_i^{\text{oversight}}$$

where:

- V_i^{funding} : Votes on CDC, BARDA, and global health security appropriations
- V_i^{IHR} : Support for International Health Regulations compliance and funding
- V_i^{surge} : Votes on pandemic surge capacity and stockpile funding
- $V_i^{\text{oversight}}$: Engagement with health security oversight

Electoral layer specification. Global health PACs commit to score-dependent independent expenditures supporting high-scorers in competitive races.

Post-office layer specification.

- Tier 1: WHO advisory positions, global health security fellowships, pandemic preparedness foundation leadership
- Tier 2: General health policy positions
- Tier 3: Pharmaceutical industry positions remain available but global health leadership positions closed

Calibration notes. Pharmaceutical industry lobbying exceeds \$300M annually. However, pandemic preparedness funding faces less concentrated opposition than defense or fossil fuels, suggesting lower IAB funding thresholds may achieve incentive compatibility.

9.9 B.4 Comparative Analysis

Table 18 summarizes key parameters across domains.

Table 18: Comparative IAB Parameters Across Domains

Domain	Opposition Cost c_i	Existing Lobbying	IAB Threshold \bar{F}
Climate	\$2–10M	\$150M+ (2024)	~\$1B
Nuclear	\$2–8M	\$130M+ (2024)	~\$500M
Pandemic	\$1–5M	\$50–150M	~\$300M
Prag- matic Clinical Trials	\$1–4M	\$20–100M	~\$200M

These estimates are illustrative and require empirical calibration. The general pattern suggests that domains with lower concentrated opposition (pandemic preparedness, health research) require lower IAB funding to achieve incentive compatibility, while domains with entrenched industrial opposition (climate, defense) require larger investments.