

A watercolor-style map of the Brazilian Amazon region. The background is composed of various shades of green, representing different types of vegetation or land use. A prominent feature is a winding road, depicted with a white center and a brown border, that meanders across the map. Several simple line-art icons are scattered throughout the map, including a tree with a rectangular canopy, a house-like structure, a square with an 'X' inside, and a square with a smaller square inside. The overall aesthetic is artistic and illustrative.

WELFARE-BASED ROAD INVESTMENT CHOICES IN THE BRAZILIAN AMAZON

Thais Vilela
Econ4Conservation



Welfare-Based Road Investment Choices in the Brazilian Amazon

AUTHOR:

THAIS VILELA, ENVIRONMENTAL ECONOMIST

Cover illustration: conceptual artwork generated via artificial intelligence, representing the intersection of infrastructure corridors and environmental conservation.

ABSTRACT

In this report, we develop a dynamic, corridor-level welfare model to identify how a planner should split a single annual budget when mobility gains, pavement decay, and forest externalities all move together over time in the Legal Amazon context. Under the assumptions used here, we find that regional budget allocations should target an 80-95% maintenance share, anchored near 95% under our main specifications. This result is robust to different parameter calibrations.

Suggested citation: Vilela, T. (2026). Welfare-Based Road Investment Choices in the Brazilian Amazon. Technical Report. Econ4Conservation. Los Angeles, United States.

Copyright © 2026 Thais Vilela / Econ4conservation. All rights reserved. The author welcomes the use of this material for non-commercial purposes, such as policy discussions, academic research, or educational activities, under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.



TABLE OF CONTENTS

List of tables.....	3
List of figures.....	3
Executive summary	4
1. Introduction.....	6
2. Literature & contribution.....	8
3. Conceptual model	9
3.1. Model dynamics.....	10
4. Formal model.....	12
5. Data and construction of the corridor panel	22
5.1. The corridor panel	22
5.2. Exogenous covariates	27
6. Results	29
7. Sensitivity and robustness.....	32
8. Limitations	35
9. Conclusion.....	37
References.....	41



LIST OF TABLES

Table 1. Calibrated model parameters	19
Table 2. Traffic index: imputation rule	25
Table 3. Data sources	28
Table 4. Scenario outcomes (central calibration, 12-year horizon and an annual discount rate of 8.5%).....	30
Table 5. Sensitivity analysis: access-benefit weight	34
Table 6. Sensitivity analysis: deforestation parameters	35

LIST OF FIGURES

Figure 1. Welfare grid.....	30
-----------------------------	----



EXECUTIVE SUMMARY

Brazil's federal highway program in the Legal Amazon faces a recurring dilemma: how to allocate a budget between maintaining the existing network and constructing new links? While project-level cost-benefit analyses and spatial models of roads and deforestation address pieces of this debate, they do not jointly answer how a planner should split a single annual budget when mobility gains, pavement decay, and forest externalities all move together over time.

Within this context, this study develops a dynamic, corridor-level welfare model encompassing 1,242 federal highway corridors across the Legal Amazon. The model integrates diverse spatial and administrative datasets.

In this model, the planner chooses a single aggregate variable each year: the maintenance share of the federal budget (M), with network expansion receiving the remainder ($1 - M$). For each candidate budget split, the model simulates 12 years of road condition, market access, deforestation pressure, and welfare (concave mobility benefits, fiscal costs with a Brazilian marginal cost of public funds, and environmental damages from carbon and non-carbon ecosystem services). The implementation finds the welfare-maximizing M by grid search under a parametric allocation rule. The model is a calibrated partial-equilibrium comparison: prices, wages, migration, and the budget envelope are exogenous.

Under the central calibration, the welfare-optimal split is about 95% maintenance and 5% new construction. This represents an interior optimum; both pure expansion and pure maintenance yield lower welfare than the optimal mix. The welfare curve is concave and flat near the peak, indicating that policy rankings are more stable than the absolute gain from small reallocations.

Why maintenance dominates (without ruling out expansion):

1. Diminishing returns to quality and access.
2. Quality-access complementarity—connectivity gains require maintained roads.
3. Hard corridor caps—spending above ceilings is fiscal waste.



Expansion-priority paths accumulate roughly 2.2 more units of total deforestation over the horizon than maintenance-priority paths.

From a policy perspective, we highlight four main recommendations:

- Adopt a maintenance-first default. Budget allocations should target an 80-95% maintenance share, anchored near 95% under our main specifications.
- Respect corridor caps. Institutional ceilings must be enforced; funding directed to corridors beyond their absorptive capacity acts as a deadweight fiscal loss to society.
- Implement protection-aware expansion. Any expansion spending should selectively favor well-enforced corridors or include explicit, funded environmental mitigation frameworks.
- Commit to annual recalibration. The model should be updated annually with fresh DNIT traffic, pavement quality, and MapBiomas deforestation data to guide rolling budget cycles.

It is worth mentioning that the model does not replace project-level EIA/RIMA or HDM appraisal; it does not pick individual projects; it does not monetize welfare as audited GDP; and it does not say maintenance always beats expansion—only that both corners are dominated and the optimum is strongly maintenance-tilted under Amazon-calibrated parameters.

Indeed, taking all this together, the model concludes that a planner who internalizes mobility, fiscal distortion, pavement dynamics, and forest externalities should devote most of a fixed federal envelope to keeping the existing network in service, while keeping a small positive share for expansion.



1. INTRODUCTION

Despite growing evidence of the Amazon forest's central role in global climate regulation (Malhi et al. 2021)¹ and biodiversity conservation (Guayasamin et al. 2021)², political and financial pressure to expand roads in the Brazilian Legal Amazon remains strong (Cezne & Otsuki 2025; Wenzel 2026). Road expansion is often justified by expected gains in market access, local production, and social welfare. Yet existing evidence suggests that the benefits of road expansion in the Amazon forest are often outweighed by its resulting environmental and social costs, particularly in pristine forest areas (Vilela et al. 2020; Araujo et al. 2023; Hanusch 2023).

As a result of this situation, the federal government faces a challenging allocation problem in which it must allocate its road infrastructure budget in a way that increases social welfare, while reducing environmental and social costs in the Amazon. Within that budget, the choice is between two competing uses: maintaining the existing network and expanding it into frontier and underserved areas. While maintenance preserves the service quality and accessibility of established corridors, expansion adds new infrastructure where access is currently limited. To our knowledge, there is currently no formal network-level framework in federal highway planning for the Legal Amazon that welfare-optimizes the split between maintaining the existing network and expanding it, while monetizing deforestation and related environmental damages in the same objective function.³ This report aims to address this gap.

We developed a calibrated dynamic welfare model of federal highway investment in the Brazilian Legal Amazon, defined over 1,242 federal corridors⁴ and parameterized

¹ As Malhi et al. (2021) observe, the Amazon is transitioning from a reliable climate regulator to a potential climate liability. This shift elevates the urgency of conservation efforts as the forest now carries the risk of becoming a significant carbon source.

² Guayasamin et al. (2021) state that "the Amazon constitutes the greatest concentration of biodiversity on Earth, with >10% of the world's described species compressed into only about 0.5% the Earth's total surface area."

³ Brazilian transport planning, such as the National Logistics Plan, tends to focus on identifying bottlenecks and infrastructure needs to reduce transport costs and potentially increase Brazilian Gross Domestic Product (GDP).

⁴ These 1,242 corridors, totaling 24,379.3 km, represent the network across nine states in the Legal Amazon for which Annual Average Daily Traffic data were available.



using publicly available data. The model integrates four variables: road deterioration, accessibility at the corridor level, deforestation responses to road condition and accessibility, and monetary valuation of carbon and biodiversity damage. Given a federal highway budget, the model returns a single number—a "maintenance share" between zero and one—at which total social welfare is maximized, together with the welfare cost of departing from that share in either direction.

Our central finding is that, under our calibrated specification, the welfare-maximizing share of the federal highway budget for maintenance in the Brazilian Legal Amazon is 95%, ranging from 76% to 98% in the sensitivity analysis.

In 2025, the *Departamento Nacional de Infraestrutura de Transporte* (DNIT) spent 74% of its national construction-and-maintenance budget on maintenance and 26% on new construction (DNIT 2026)⁵—already in the maintenance-favorable range our model identifies, though slightly below the optimum. In the Brazilian Legal Amazon, additional federal highway spending generates more social welfare when used to maintain existing roads than when used to build new ones. This is consistent with the general-equilibrium evidence on the deforestation costs of road expansion in the Amazon (Araujo et al. 2023, 2025).

The remainder of this report is organized as follows. Section 2 reviews the empirical literature and positions our contribution. Section 3 introduces the model conceptually, while Section 4 sets out the formal mathematical framework. Section 5 describes the data and the construction of the corridor panel used for calibration. Section 6 presents our baseline results and welfare-maximizing budget allocations. Section 7 evaluates the stability of these findings through a series of sensitivity analyses, and Section 8 highlights the structural limitations of the model along with their potential implications for our results. Finally, Section 9 concludes with an evaluation of policy implications and highlights some directions for future research.

⁵ DNIT (2026) reports a total road infrastructure investment of R\$ 12.91 billion in 2025. This budget was distributed among maintenance (R\$ 9.26 billion), new construction (R\$ 3.24 billion), and operations (R\$ 0.41 billion). The investment ratio between maintenance and expansion stands at 74% and 26%, respectively, when operations are excluded.



2. LITERATURE & CONTRIBUTION

Research on the welfare effects of transport infrastructure has historically emphasized its benefits. Indeed, a large literature in spatial economics documents positive welfare effects of transport network expansion, operating through reduced trade costs, increased market integration, and interregional price convergence (Allen & Arkolakis 2014; Donaldson 2018; Pellegrina & Sotelo 2025). However, as these studies are often conducted in contexts where infrastructure is already established or environmental constraints are secondary, they typically do not internalize the environmental externalities associated with new construction.

This omission is particularly problematic in the Amazon, where the impact of the transport network on natural resources is of first-order importance. The literature on the environmental costs of infrastructure in tropical biomes is large and converges on a clear finding: transport expansion in the Amazon accelerates deforestation, often resulting in social and environmental costs that outweigh the financial benefits of the projects (Pfaff 1999; Soares-Filho et al. 2006; Vilela et al. 2020; Souza & Silva 2022; Zhou et al. 2026).

More specifically, the literature identifies two main channels through which transport infrastructure drives forest loss. First, a direct channel, whereby the physical footprint of construction requires the immediate clearing of forest along the roadbed and its adjacent buffers (Laurance et al. 2009; Lapola et al. 2023). Second, an access channel, through which infrastructure lowers the cost of subsequent forest conversion for agriculture, logging, and extractive activity (Pfaff 1999; Asher et al. 2020).

Concerns about the deforestation effects of Amazon roads date to descriptive and remote-sensing studies of the 1980s and early 1990s (Fearnside 1987; Skole & Tucker 1993). Peer-reviewed reduced-form evidence has documented these effects since at least Pfaff (1999), with subsequent work refining the importance of road hierarchy (Asher et al. 2020). For example, in the Amazonian context, primary arteries amplify forest loss through a "snowball effect": official road expansion acts as a catalyst for the proliferation of unofficial and illegal secondary networks, which are responsible for



much of the deforestation footprint (Barber et al. 2014). Simulations of the BR-163 paving corridor echo this dynamic, projecting that the primary paving would unlock a decadal buffer of secondary clearing (Soares-Filho et al. 2006).

More recent quantitative spatial general-equilibrium analyses confirm and sharpen these estimates by accounting for spatial spillovers, in which improved market access across the transport network triggers deforestation in regions outside the immediate project corridor. Araujo et al. (2025) report that a one percent increase in market access raises deforestation by approximately 0.5 percent in the Brazilian Amazon—an effect that is robust across different estimation strategies, alternative trade-elasticity calibrations, and alternative transport-cost parameterizations.

Taken together, this literature has documented the magnitude of the environmental cost of road infrastructure in the Amazon with growing precision. From a planner's perspective, however, a gap remains between what the literature provides and what is needed to allocate a fixed federal budget between maintenance of existing roads and expansion through new construction. This report addresses that gap by developing a calibrated model that integrates environmental costs alongside transport welfare gains within a single planner's decision framework. The central question the model aims to answer is the following: what is the welfare-optimal split of Brazil's federal highway budget in the Legal Amazon between maintenance of existing roads and expansion through new construction, once environmental and social externalities are internalized?

3. CONCEPTUAL MODEL

We develop a dynamic decision model in which the planner decides each period how to split a limited federal highway budget across two types of spending in the Brazilian Legal Amazon: maintenance of existing roads and network expansion, knowing that today's choices affect future road quality, connectivity, and deforestation.

As with most economic models, we assume that the planner's objective is, in each period, to maximize social welfare which, in this study, has three components: transport



benefits, fiscal (i.e., government budget) costs, and environmental costs. More specifically:

- Transport benefits: better road condition and enhanced connectivity reduce travel times, lower transportation costs, and improve market access, thereby stimulating regional economic activity (Herrera Dappe et al. 2024). While these gains increase as the network improves, they exhibit diminishing marginal returns, meaning that each incremental upgrade yields progressively smaller economic benefits. Furthermore, realized access gains depend on the baseline quality of the road network, reflecting a vital quality-access complementarity where expansion is only as effective as the maintenance supporting it. The model aggregates these localized dynamics across the entire network by weighting each corridor according to the size of the population it serves.
- Fiscal costs: these costs represent the direct public spending required for both network maintenance and expansion. In this study, we assume that the total transport budget is fixed each period; the planner chooses only its split between maintenance and expansion. To reflect absorption limits, the model applies hard caps at the corridor level, i.e., only spending up to each cap improves road condition or access. If caps prevent the full budget from being productively used, that unproductive share is still counted as a fiscal cost, while generating no additional network improvements.
- Environmental costs: in this study, we treat deforestation as a social cost of road investment. Deforestation rises with connectivity and new construction and falls with protection. We assume that expansion is the stronger environmental channel, while maintenance has smaller associated forest impacts. In our model, monetized deforestation damages are subtracted from welfare.

In sum, in each period, welfare is equal to the benefits from the network, minus the fiscal costs and the environmental costs. The government maximizes the discounted sum of this welfare over the planning horizon, so that future periods count but less than the present.

3.1. MODEL DYNAMICS



To account for the intertemporal trade-offs between maintenance and expansion, we develop a dynamic model where current budget decisions shape long-term outcomes. In this setting, maintenance spending preserves serviceability on the existing network, preventing road deterioration and the erosion of past capital investments. Conversely, expansion spending extends connectivity into new frontiers, triggering environmental consequences that unfold over time. A static analysis would miss these temporal adjustments, failing to capture how today's budget split shapes future road quality, market access, and deforestation pressures. In this study, we assume that this underlying dynamism operates through three interconnected mechanisms:

- Road condition: this reflects how well existing roads are functioning. Usually, the condition of a road worsens over time because of use and weather, and it worsens faster where traffic is heavier. Maintenance spending slows or reverses this decline. To capture this situation, the model has the following dynamic: if the government neglects maintenance, condition falls; if they invest in maintenance, condition improves or stabilizes.
- Connectivity: access and connectivity reflect how well places are linked to markets and to each other. Connectivity improves when the planner spends on expansion and can also respond to improvements or deterioration in road condition. The model considers that expansion raises long-run connectivity while maintenance preserves the quality of existing links and thus supports the level of access that the current network can deliver.
- Deforestation: forest loss in each period depends on market access, connectivity, road condition, and current-period expansion expenditures. Areas with stronger environmental protection or institutional enforcement are assumed to have lower deforestation rates for any given level of access. In this framework, maintenance affects forest pressure indirectly by preserving infrastructure quality along established routes (the intensive margin), whereas network expansion drives deforestation directly through new-construction spending that pushes into pristine frontiers (the extensive margin).

These three dynamics link today's spending to tomorrow's benefits and costs. When solving the optimization problem, the planner accounts for the fact that underinvesting



in maintenance today leads to worse condition and possibly lower benefits later, while overinvesting in expansion can result in higher deforestation and environmental costs for the future.

In terms of constraints, the planner faces two main constraints:

- Budget: Each period, the planner faces a fixed federal budget. This is split into a maintenance pool and an expansion pool. Corridor-level spending is allocated within each pool (subject to caps), not from a single undifferentiated pool.
- Feasibility: spending in each corridor is bound above (and below by zero). These bounds reflect physical and institutional limits (real-world bottlenecks) on how much can be done in each place in each period.

4. FORMAL MODEL

In this study, we assume the following planner's problem:

$$\max_{\{m_{i,t}, n_{i,t}\}} \sum_{t=0}^T \beta^t \cdot W_t$$

Where:

t represents time (year) and i road corridor.

$m_{i,t}, n_{i,t}$ are the control variables corresponding to the budget share that will be allocated to maintenance and new road construction, respectively.

$0 < \beta < 1$ is the time preference defined by the discount rate (d) : $\beta = \frac{1}{(1+d)}$

and $W_t = B_t - C_t - E_t$ for every t where:

- B_t : Benefits related to connectivity and market access
- C_t : Fiscal costs related to road maintenance and new road construction
- E_t : Environmental costs caused by roads

Benefits

For corridor i and year t , the transport benefit equation is the following:

$$B_{i,t} = w_i \cdot [b_q \cdot q_{i,t}^{\gamma_q} + b_a \cdot a_{i,t}^{\gamma_a} \cdot q_{i,t}^{\beta_{qa}} + b_p \cdot pop_i]$$

where:

- w_i : Equity weight
- $q_{i,t}$: Road condition index
- $a_{i,t}$: Access/connectivity index
- pop_i : Population

As described in the conceptual framework, the benefits in our model comprise three components: pavement quality (i.e., road condition), network connectivity, and population served. To incorporate distributional concerns, the benefits are scaled by an equity weight, ensuring that infrastructure improvements yield higher social welfare when they occur in more poor states. It is worth mentioning that this weight does not constitute a full distributional welfare analysis, as it abstracts from within-state inequality:

$$w_i = 1 + poverty\ index_i$$

In terms of functional structure, both road condition and connectivity enter the benefit equation concavely to ensure an interior solution. The exponents $\gamma_q, \gamma_a \in (0,1]$ capture the diminishing marginal returns and are set to $\gamma_q = \gamma_a = 0.7$ to impose moderate diminishing returns on road condition and access. It is worth mentioning that this value is not taken from an empirical estimate; it lies in the interior of a plausible range consistent with concave benefit functions in transport and trade applications (Donaldson, 2018; Allen and Arkolakis, 2014).

In the case of network connectivity, beyond experiencing diminishing returns, it also exhibits a fundamental complementarity with road condition. This interaction is captured by the parameter $\beta_{qa} \in [0,1]$. When $\beta_{qa} = 0$, then access benefits do not depend on quality, but as β_{qa} approaches 1, connectivity gains require maintained roads. In this study, we assume a moderate interaction of $\beta_{qa} = 0.5$, implying that while

poor road conditions degrade transport benefits, they do not reduce access completely.

Aggregating these dynamics, total period benefits are defined as: $B_t = \sum_i B_{i,t}$

Fiscal costs

To keep the model simple, we assume fiscal costs are linear in public spending. The constant shadow prices c_m and c_n represent the marginal cost of public funds, capturing the distortionary welfare cost of taxation (Lanzer 2011).⁶ Under hard corridor capacity caps⁷, the total fiscal cost in period t is expressed as:

$$C_t = c_m \cdot m_t^{budget} + c_n \cdot n_t^{budget} \text{ and } m_t^{budget} + n_t^{budget} = \text{Transport budget (fixed)}$$

where m_t^{budget} and n_t^{budget} are the aggregate maintenance and expansion budget pool, which sums to a fixed annual funding budget. Crucially, when the marginal cost of public funds is symmetric across spending types ($c_m = c_n$) (which is what we assume here), the total societal fiscal cost C_t remains invariant to how the budget is split. Consequently, the planner's optimization problem does not revolve around minimizing total fiscal expenditures; instead, the choice comes down to balancing (a) the economic benefits of the roads, (b) the environmental damage from deforestation, and (c) the risk of wasting money if a project receives more funding than it can actually handle.

Environmental costs

In the case of environmental costs, we assume they are entirely driven by deforestation. More specifically, for a corridor i in period t , the damages are split into carbon and non-carbon components:

$$E_t = \sum_{i=0}^N (w_{carbon_t} \cdot e_{i,t} + w_{bio} \cdot e_{i,t})$$

⁶ Each R\$ 1 spent by the planner costs society more than just R\$ 1 because, to raise those funds, the government taxes its citizens or borrows money, creating distortions in the economy.

⁷ Hard corridor capacity caps mean that each corridor can absorb only a limited amount of maintenance and expansion spending per year; spending above a corridor's ceiling cannot be used productively and, under our hard-cap rule, cannot be reallocated to other corridors.



where w_{carbon_t} is the climate-related (carbon) impact weight, w_{bio} represents the non-carbon damage weight, and $e_{i,t}$ is a deforestation flow index. The latter captures how a corridor's baseline deforestation interacts with infrastructure-related pressures and protection. It is described as:

$$e_{i,t} = \max\{0, \text{baseline } e_i + \eta_a a_{i,t} + \eta_q q_{i,t} + \eta_{new} n_{i,t} - \eta_{prot} prot_i\}$$

where:

- baseline* e_i : Baseline deforestation index (MapBiomass, 10 km buffer).
- $a_{i,t}$: Road connectivity / market access.
- $q_{i,t}$: Road condition
- $n_{i,t}$: Expansion allocation.
- $prot_i$: Protection (protected area overlap x enforcement).

The specification above captures baseline heterogeneity, infrastructure channels, and linear mitigation.

Baseline heterogeneity: each corridor has a fixed baseline deforestation index (*baseline* e_i) which enters additively and sets how much deforestation pressure the corridor starts with in every period.

Infrastructure channels: the planner affects deforestation through two routes. Expansion spending adds clearing pressure directly in the same period. Maintenance spending has no direct term in the deforestation equation; it matters indirectly by improving road condition and market access over time, and those states enter the clearing index with coefficients η_a and η_q .

Mitigation: conservation overlap and enforcement enter as a fixed subtractive shield ($-\eta_{prot} prot_i$). Corridors with stronger protection have lower clearing flows, all else equal. Because mitigation is linear and additive, it can offset high baseline pressure in levels, but it does not remove the additional clearing triggered by a given road upgrade.



It is worth noting that because the baseline deforestation index reflects historical clearing near each corridor (from 2014 to 2024), it should be interpreted jointly with the η parameters. Together, they function as a reduced-form deforestation block, not an isolated causal estimate of road-driven deforestation.

Since our model is dynamic, we define two transition equations.

(a) Road condition dynamics

In this study, we assume that while road condition decays over time through natural weathering and traffic stress, it is restored through maintenance spending. We capture this with a standard transition function in which next period's quality equals current quality, minus a fixed share lost to depreciation and traffic-related wear, plus the improvement from maintenance investment in that period:

$$q_{i,t+1} = \min(\max(q_{i,t} \cdot (1 - \delta_q - \varphi_{traffic} \cdot traffic_i) + \alpha_m \cdot m_{i,t}, 0), 1)$$

Where:

- $q_{i,t}$: This is the normalized road-condition index for corridor i at t (where 0 represents worst condition and 1 is perfect condition).
- $q_{i,t+1}$: This is road condition in the next period (after the government makes its maintenance decisions and after depreciation).
- δ_q : This represents the fixed percentage the road degrades each period due to weather and aging, regardless of traffic.
- $\varphi_{traffic}$: This determines how severely traffic volume accelerates degradation.
- α_m : This parameter converts every dollar (or unit) of maintenance spending into an improvement in road quality (i.e., converts maintenance allocation into condition improvement).
- $m_{i,t}$: The maintenance budget allocated to corridor i in period t .

Note that this function enforces dynamic constraints which ensure that you cannot maintain a brand-new road to make it "more than new" (clamping at 1), and that if you neglect a road with heavy traffic, it cannot fall below zero.

(b) Connectivity dynamics



Similarly to road condition, connectivity evolves over time. Next period's connectivity equals current accessibility (or connectivity), plus gains from expansion spending, plus the effect of recent changes in road quality:

$$a_{i,t+1} = \min(\max(a_{i,t} + \alpha_n \cdot n_{i,t} + \alpha_q \cdot (q_{i,t} - q_{i,t-1}), 0), 1.5)$$

Where:

- $a_{i,t}$: This is the state of road connectivity for corridor i at period t .
- $a_{i,t+1}$: This variable represents next period road's connectivity.
- $n_{i,t}$: This represents the new construction or capacity expansion. The term $\alpha_n \cdot n_{i,t}$ adds capacity and structural connectivity to the system. So, new construction improves long-run connectivity.
- $q_{i,t} - q_{i,t-1}$: This is the change in road quality (or condition) from the last period to this period. If the government maintained the road in a way that $q_{i,t} > q_{i,t-1}$, then accessibility increases. Otherwise, quality collapsed, resulting in decreased realized accessibility.
- α_q : It determines how much a sudden drop (or rise) in road condition impacts the realized speed, safety, and reliability (and thus the accessibility) of the corridor.

Note that this function is also clamped. We assume that the road cannot have negative accessibility. If a road completely closes due to neglect or construction collapse, accessibility is 0. On the other hand, we clamp the function at 1.5. This implies that a corridor's connectivity cannot exceed 1.5 on the normalized index. In other words, we assume that even with infinite budget and effort applied to both expansion and maintenance, road connectivity cannot grow without limit on the normalized scale.

(c) Deforestation equation

In this simplified model, we assume that deforestation is not a state variable in the sense that its current state does not evolve from period to period. In our case, deforestation in each period is an outcome determined by that period's access, road condition, and expansion spending; it is not a stock that carries over from one period



to the next in the model's equations. Based on this, the deforestation equation is the following:

$$e_{i,t} = \max(0, base_{e_i} + \eta_a \cdot a_{i,t} + \eta_q \cdot q_{i,t} + \eta_{new} \cdot n_{i,t} - \eta_{prot} \cdot Prot_i)$$

Where:

- $e_{i,t}$: This is the normalized deforestation index.
- $base_{e_i}$: This variable corresponds to the existing deforestation in corridor i . It means that even when the other terms are zero (or at their reference levels), there can still be some background deforestation in that corridor.
- $Prot_i$: This variable represents the protection level. It can also be interpreted as the mitigation factor. It might represent, for example, government presence, environmental regulations, or conservation units that act as a buffer against deforestation. However, it is worth noting that, in this study, protection is a corridor characteristic, not a policy lever that changes within the simulation.

With this equation, we aim to capture the trade-offs described initially in this report. In our model, when the planner decides to spend money on expansion or maintenance to maximize welfare, it can effectively increase deforestation. More specifically, by increasing road quality and accessibility, the planner satisfies the goal of economic connectivity, but it also adds forest pressure (η_a and η_q) to the index. The same is true in the case of new roads.

The use of $\max(0, \dots)$ is an economic constraint. It prevents the model from suggesting that a road with high protection and zero maintenance would create forest (or negative deforestation).

Parameter calibration

Table 1 summarizes the parameters used in the model. It is worth noting that the structural parameters in this model are literature-informed, not literature-estimated. Where external work reports elasticities, damage valuations, or engineering ranges, we

use those results to fix signs, relative magnitudes, and plausible bands. Where the model works with normalized corridor indices such as Hansen-style access, pavement-quality scores, and MapBiomass-based deforestation flows, coefficients such as η_a are normalization choices, i. e., they set how strongly each index enters the equations while keeping outcomes inside defensible ranges and consistent with internal cross-checks.

TABLE 1. CALIBRATED MODEL PARAMETERS

Block	Parameters
Deforestation	$\eta_a = 0.35, \eta_q = 0.05, \eta_{new} = 0.5$, and $\eta_{prot} = 0.3$
Dynamics	$\delta_q = 0.07, \varphi_{traffic} = 0.05, \alpha_m = 0.8, \alpha_n = 0.8, \alpha_q = 0.10, q_{max} = 1$, and $a_{max} = 1.5$
Environmental damage	$scc = 1, c_{bio} = 1, growth = 0.015$ per year
Benefit weights & MCPF	$b_q = 1, b_a = 1.5, b_{pop} = 0.4$, and $c_m = c_n = 1.17$
Curvature and complementarity	$\gamma_q = \gamma_a = 0.7, \beta_{qa} = 0.5$, and hard caps = true

We do not treat any such coefficient as if it were read directly from a published table unless the mapping is explicit (e.g., the social discount rate and the marginal cost of public funds). Sensitivity analysis (Section 7) varies the parameters that move the optimal maintenance budget share m^* the most: quality-access complementarity (β_{qa}), the access-benefit weight (b_a), the environmental-damage scale, and the discount rate.

Fiscal flows and the federal budget are expressed in R\$ billion per year (bn/year) on the Legal Amazon federal-highway panel (1,242 corridors). Benefit weights b_q, b_a , and b_{pop} are expressed relative to $b_q = 1$, which serves as the welfare numéraire. The same is true for environmental damages which is expressed in comparable welfare units.

Constraints

To fully describe the planner's problem, we must now account for the model constraints. In this simplified model, we have two sets of constraints: one related to budget and another one related to implementation bounds.



Regarding the budget, we assume that, in each period, the planner faces a fixed federal envelope, G_t , which is split into a maintenance pool and an expansion pool:

$$m_t^{budget} = m \cdot G_t, \text{ and } n_t^{budget} = (1 - m) \cdot G_t$$

where m represents the maintenance share.

Additionally, for both $m_{i,t}$ and $n_{i,t}$ there are resource constraints⁸:

$$0 \leq m_{i,t} \leq m_i^{max}$$

$$0 \leq n_{i,t} \leq n_i^{max}$$

Where:

m_i^{max} : the maximum amount of resource that can be allocated to road maintenance.

n_i^{max} : the maximum amount of resource that can be allocated to constructing new roads.

Now that we have fully described the model, we note that there is not a close-form solution for $m_{i,t}$ and $n_{i,t}$ mainly for the following:

1. Coupling: the problem links all corridors within the same budget constraint in each period; thus, the optimal investment in one corridor depends on the state of all others and the marginal value of states over time.
2. Non-smoothness: the transition and deforestation equations include non-smooth operations (e.g., upper/lower bounds on states and non-negativity constraints on environmental flows). These create distinct regimes where the valid optimality conditions change.
3. Temporal recursivity: even where fiscal costs are linear in the budget pools, transport benefits are concave in road condition and access, and the dynamics link today's budget split to future road condition and access and deforestation.

⁸ In future exercises, additional constraints, such as minimum maintenance floors for strategic corridors, could be added to the model to make it more complex.



The decision at time t therefore depends on the whole future path and does not reduce to a simple closed-form rule for corridor-by-corridor spending.

In the literature, this challenge is addressed through three compatible approaches:

- Characterizing necessary conditions: using optimality conditions (e.g., the Bellman system with co-states, or Karush-Kuhn-Tucker conditions when reformulated as a non-linear program);
- Numerical methods: obtaining the optimal policy through dynamic programming or convex/non-convex optimization over the control vector;
- Parameterized policy rules: restricting the policy to an explicit allocation rule and selecting parameters through welfare maximization or scenario comparison.

Any of these approaches preserve the economic substance of the problem—the trade-offs between mobility benefits, fiscal costs, and environmental damages—without requiring $m_{i,t}$ and $n_{i,t}$ to appear as closed functions of a few observable variables.

Thus, the operational solution, whether derived via numerical optimality conditions or a calibrated parameterized rule, constitutes the empirical implementation of the decision program, rather than an arbitrary substitute for the theoretical problem.

Implementation strategy: heuristic allocation via policy optimization

In the current implementation, the complexity of the general optimization problem is addressed through a simplified Parametric Allocation Rule. Instead of solving for independent investment levels for every corridor in every period, the model operates as follows:

Step 1. For each period, a global allocation parameter is established, where a fixed fraction of the total budget is dedicated to maintenance (m) and the remainder to expansion (n).

Step 2. Within each pool, funds are distributed across corridors in proportion to need scores – maintenance need rises with quality gaps, traffic, and equity weight;



expansion need rises with access gaps, traffic, and population exposure – and then clipped at corridor ceilings (m_i^{max} and n_i^{max}).

Step 3. The system executes a numerical grid search across the full spectrum of the allocation parameter (from 0% to \$100% maintenance). It simulates the long-term trajectory of the system for each possible "mix."

Step 4. The model identifies the specific fraction that maximizes the discounted aggregate social welfare over the entire planning horizon.

Therefore, rather than attempting to solve the general high-dimensional optimization, the implementation solves for the optimal policy mix within a structured ruleset. This approach allows for a robust comparison of strategic scenarios while maintaining computational feasibility.

5. DATA AND CONSTRUCTION OF THE CORRIDOR PANEL

This section describes the corridor-level panel that the model takes as input: 1,242 federal corridors in the Brazilian Legal Amazon. Following the standard decomposition used in dynamic-optimization models, the data are organized into three groups: state variables, which define the conditions each corridor faces at the start of the planning horizon, exogenous covariates, and choice variables, which are the set of variables that the planner picks each period subject to the budget and caps. In this section, we focus on the description of the first two groups.

5.1. THE CORRIDOR PANEL

The unit of analysis is the corridor as defined by DNIT in its planning, budgeting, and contracting frameworks. The analysis is restricted to the Legal Amazon, encompassing nine States: Acre, Amapá, Amazonas, Maranhão, Mato Grosso, Pará, Rondônia, Roraima, and Tocantins. Additionally, the panel includes only operational corridors: the three *Sistema Nacional de Viação* (SNV) statuses indicating segments that have not yet been built or are still under primary construction (planned, under construction, or

under paving work) were excluded, because the model's traffic-based benefits calculation has no defensible imputation peer for corridors that have never carried traffic. The remaining SNV statuses – unpaved/dirt, urban crossing, under expansion, paved, and multi-lane – were retained. As a result of these filters, we ended up with a sample of 1,242 corridors, spanning 24,379.3 km.

For each one of the federal corridors, we collect the following state variables (data):

- (a) Length and surface status. These data are collected directly from DNIT. The length corresponds to the length of the corridor, and its unit is km. The surface status corresponds to whether the road is: paved, unpaved, planned, under construction, under paving work, urban crossing, under expansion, or multi-lane.
- (b) Baseline quality or condition. This variable corresponds to the initial road condition. In this report, we assume a simplistic approach in which we translate each one of the “surface status” category into one number between 0 and 1, adding between 0.10 on the 0-to-1 scale. We deliberately keep this scale linear because the model's welfare formula already captures two non-linear effects: diminishing returns to quality, and quality-access complementarity (Section 4).
- (c) Baseline access or connectivity. Following the gravity-style market-access construction described in Araujo et al. (2022), we construct a corridor-level access index, using population data from IBGE (2025) and Haversine distance as a proxy for transport cost.⁹ For each corridor i , access is initially calculated as:

$$Access_i = \sum_{j=1}^J \frac{Pop_j}{\max(d_{ij}, d_{min})^\beta}$$

Where j corresponds to the approximately 5,570 Brazilian municipalities, Pop_j is the population in municipality j , d_{ij} is the distance between corridor i 's midpoint and the municipality j 's centroid, and d_{min} is assumed to be 5 km to prevent the term $\frac{Pop_j}{d_{ij}^\beta}$ from exploding in the case of municipalities sitting on top

⁹ Araujo et al. (2022) calculate transport cost using freight data.

of a corridor midpoint, and β is a distance decay parameter, i.e., it is a parameter that gives greater weight to nearby population than to distant population. Without it, the model would assume that every person in the country, no matter how far away, has an equal chance of using that specific road. In this report, we assume β to be 1.5.¹⁰

Once we calculate access for all corridors in our sample, we apply a log transformation and a min-max normalization to create the index, with mean equal to 0.30 and median close to 0.31. That concentration in the lower part of the [0, 1] range is the expected pattern for a population-weighted access potential on a low-density network. Qualitatively, this aligns with Araujo et al. (2022), who report that market access for a typical municipality in the Legal Amazon is about 2.1 times lower than in the rest of the country. Although our corridor-level index is not directly comparable to their municipality-level index, it exhibits the same broad spatial ordering, i.e., peripherally embedded corridors score lower than a hypothetical dense “central Brazil” sample would under the same construction.

- (d) Traffic. To create a traffic index, we follow a three-step approach. First, we calculate a corridor-level traffic volume using the Annual Average Daily Traffic (AADT) from DNIT. This step is needed because a single federal corridor can be made of several shorter segments, each with its own length and, when available, its own ADDT. Thus, we summarize all of that into one traffic value per corridor i by taking a length-weighted average:

$$v_i = \frac{\sum_s v_s L_s}{\sum_s L_s}$$

Where v_i is ADDT of the sum over the segments s , v_s is ADDT on segment s , and L_s is the segment’s length. Second, we rescale the corridor-level values to the interval [0,1] using min-max normalization across the corridors in our sample. After this step, the variable tells us how “busy” a corridor is relative to

¹⁰ As far as we know there is not a universally accepted β range that maps onto our distance-km specification. We set $\beta = 1.5$ as a mid-range gravity-style choice and verify robustness by computing rankings at $\beta \in \{1, 2\}$.



others in the Legal Amazon panel, not the absolute number of vehicles per day. Finally, in the third step, we fill in any gaps in a simple and conservative way by assigning the median traffic index among measured corridors of the same road status (Table 2).

TABLE 2. TRAFFIC INDEX: IMPUTATION RULE

Status	Missing	Filled with
Paved	208	own median (n=625 measured peers)
Multilane	83	own median (n=132)
Under expansion	5	own median (n=5, small sample but operational paved-being-duplicated)
Urban crossing	10	own median (n=8)
Unpaved /dirt	168	'Urban crossing' median (the closest lower-quality category with measured peers)
Total imputed	474	(\approx 38% of the 1,242 kept corridors)

Note: We use the median instead of the mean to prevent a small number of high-traffic corridors from skewing the results.

- (e) Population exposure. This variable summarizes how much resident population the corridor "passes through," in a simple way that respects municipal boundaries. We intersect each corridor with the IBGE municipal map. Then for every segment s where the corridor lies inside municipality $m(s)$, we calculate the length L_s of that segment and the municipality's population $P_{m(s)}$. As with the traffic index, the length-weighted mean population along corridor i is:

$$\text{Length weighted mean pop}_i = \frac{\sum_{s \in i} L_s \cdot P_{m(s)}}{\sum_{s \in i} L_s}$$

Where $\sum_{s \in i} L_s$ adds over all intersection segments on corridor i . For example, suppose a corridor (road) crosses 3 municipalities, then we will have 3 segments, lengths L_1, L_2, L_3 and populations P_1, P_2, P_3 . Based on this, we calculated the length-weighted mean population along the road as:

$$\text{Length weighted mean pop}_i = \frac{L_1 \cdot P_1 + L_2 \cdot P_2 + L_3 \cdot P_3}{L_1 + L_2 + L_3}$$



- Population is held constant within each municipality.¹¹ Finally, we min-max normalize across all corridors in the Legal Amazon panel to obtain an index in [0, 1], comparable to the other corridor scores. As with traffic, this is a relative measure: it ranks corridors by population exposure within the panel, not by absolute headcounts.
- (f) Protection. The protection index is the share of the corridor's length that lies inside federal protected areas (PA)—data collected from ICMBio.¹² The share is calculated using the formula: $\frac{\text{Length inside PA}}{\text{Total length}}$. We note that corridors with no overlap receive zero.
- (g) Baseline deforestation. This variable corresponds to a single number per corridor that acts as a scale or intercept for the deforestation block, so that corridors with different recent land-use histories do not enter the model as if they were identical. We use MapBiomass annual land-use / land-cover (Collection 10.1). We specifically track "primary vegetation suppression" (MapBiomass Code 4) to ensure the analysis focuses on the conversion of native forests rather than transitions into secondary vegetation or already altered land. Then, we calculate how many times each pixel was classified as cleared over a 11-year period (2014-2024). Rather than measuring simple hectares, this identifies areas with persistent or repeated clearing exposure near the infrastructure. Finally, we apply a 10 km buffer around each highway corridor and calculate the average clearing frequency within that zone.¹³ The final value is rescaled to a 0-1 index, representing the average annual probability of clearing occurrence, ensuring the deforestation variable is consistent with other indices used in the dynamic model.

¹¹ We acknowledge the limitation of using total municipal population as uniform within each municipality's polygon since rural stretches inside populous municipalities can receive more population weight than they would under a finely disaggregated (e.g., gridded) population map, and urban cores are not weighted more heavily than remote parts of the same municipality.

¹² Non-federal protected areas and indigenous areas are not considered in this report.

¹³ The buffer is not a claim that all clearing within 10 km is caused by that exact federal segment. However, it is a standard operational neighborhood for "road-adjacent pressure" in Amazon land-use work following seminal studies such as Barber et al. (2014), and Botelho et al. (2022).



(h) Maintenance and new construction caps. These two variables correspond to the per-corridor caps (i.e., the maximum amount of resources the planner can allocate to each one). They are calculated using the percentages of DNIT's spending presented in *Boletim da Infraestrutura* (DNIT 2025a). According to the report, maintenance and conservation together accounts for 40.412% (M_{share}) of the relevant DNIT's roads spending mix, while new construction accounts for 13.883% (N_{share}). The remaining of the budget is used for operational activities. To calculate the per-corridor cap, we use the following formula:

$$\begin{aligned} \text{Maintenance cap}_i &= M_{share} \cdot s_i \cdot \text{Annual total} \\ \text{Construction cap}_i &= N_{share} \cdot (1 - \text{protection index}_i) \cdot s_i \cdot \text{Annual total} \end{aligned}$$

Where $s_i = \frac{L_i}{\sum_j L_j}$ is corridor i 's share of total operational federal length in the Legal Amazon panel (so $\sum_i s_i = 1$), and protection index, as described above, is the overlap penalty on new construction only. The 'annual total' is the regional annual budget available for maintenance and new construction of roads assigned to the Legal Amazon (our panel) before the planner chooses the maintenance share.

Data from DNIT (DNIT 2026) suggests that, in 2025, a total of R\$ 3.24 billion was spent on road construction and upgrades, R\$ 9.26 billion to maintenance and conservation, and R\$ 411.23 million to operations; so roughly R\$12.9 billion for all roads spending and about 12.5 billion per year for maintenance and construction.

Under the assumption that 42.8% of the total national length under DNIT's management are located in the Legal Amazon (24,379 km/57,000 km), we calculate a regional annual budget of R\$ 5.35 billion per year, acknowledging that this value might overestimate the amount available for the region.

5.2. EXOGENOUS COVARIATES



In addition to the dynamic state variables, we also have three covariates which are not chosen by the planner and do not follow a transition equation such as the state variables. These covariates enter as state-level indices and are meant to summarize economic structure, socioeconomic stress, and federal environmental enforcement, i.e. the context in which road budgets are spent.

- Economic activity. We use state-level gross value added (*IBGE Contas Regionais, 2023*) divided by IBGE population estimates (2024) to build a value-added-per-capita index. We note that, in this study, we opt for using gross value added instead of gross product which double-count intermediates.
- Socioeconomic stress (poverty index). We use state-level Human Development Index (HDI) from the PNUD/IPEA/FJP *Atlas do Desenvolvimento Humano no Brasil (2022)*, and define a deprivation proxy as $1 - \text{HDI}$, then re-scale it to $[0, 1]$ by min-max across the nine states in the Legal Amazon panel. Higher values indicate lower human development (income, education, and longevity combined) compared to other states in the panel.
- Institutional oversight (federal enforcement). We construct an index from IBAMA *Autuações Ambientais (2025)*, restricted to flora infractions, counting notices by the state where the infraction occurred, and scaling by state area (infraction notices per $1,000 \text{ km}^2$), then min-max across the nine panel states. The approach is motivated by Assunção et al. (2023) on monitoring and law enforcement in the Amazon, with two caveats: (a) our measure is a state-level intensity proxy and is not comparable to their empirical design or identification strategy; and (b) state environmental agencies also enforce in some states, so IBAMA-only counts likely understate total enforcement effort.

Table 3 presents in detail the sources used to create each one of the indexes described above.

TABLE 3. DATA SOURCES

Variable	Source	Data temporal coverage
Legal Amazon boundaries	IBGE	2024

Corridor ID; state; surface (e.g., paved or unpaved); length (km)	DNIT	July 2025
Annual Average Daily Traffic	DNIT	2022
Municipalities' population	IBGE SIDRA (Table 6579)	2025
Protected areas	ICMBio	April 2026
Maintenance cap	DNIT	2025
Annual deforestation	MapBiomas Coleção 10.1	2014-2024
Gross value added	IBGE <i>Contas Regionais</i>	2023
Human development index	PNUD/IPEA/FJP	2021
Environmental infraction notices	IBAMA (forest and native-vegetation infractions)	2025

6. RESULTS

Given the calibrated planner's problem, welfare is maximized at an interior maintenance share of $m^* = 0.95$, i.e., under the model's assumptions, we find that 95% of the budget is allocated to maintenance and 5% to new construction. This optimum is interior in the sense that it is not pinned at the boundary values (both corners are dominated). Discounted total welfare rises from 13,570 at $m = 0$ to 13,640 at $m^* = 0.95$, while $m = 1$ delivers 13,639, slightly below the peak (Figure 1).

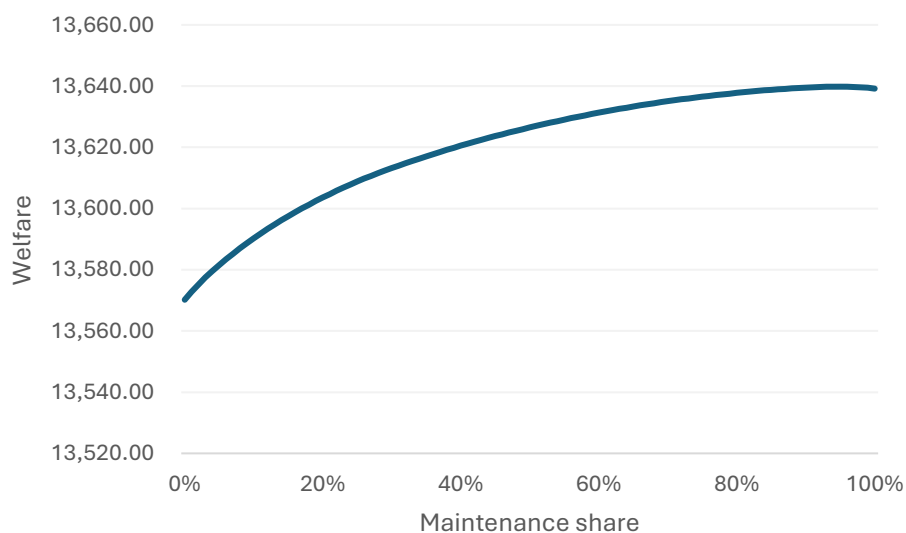




FIGURE 1. WELFARE GRID

Note: Welfare is almost flat above 90% maintenance; the optimum at 95% is interior but the welfare cost of 100% maintenance is small (approximately 0.6 units, <0.01% of present-value welfare).

The model does not recommend full maintenance; instead, it recommends a high maintenance share with a small expansion residual to capture access gains where they are not undone by environmental costs and cap constraints.

At the same time, the distance between $m = 1$ and $m^* = 0.95$ in welfare space is modest relative to the distance from low-maintenance rules to the peak. In a policy context, the central message is therefore best summarized as a maintenance-dominant allocation with non-trivial but limited room for expansion, rather than as a knife-edge sensitivity to being exactly at 95% rather than 100%.

Maintenance-priority vs expansion-priority vs welfare-optimal

To connect the optimum to simpler rules that might appear in planning discussions, the model also evaluates three fixed- m benchmarks in addition to the optimised scenario:

- Maintenance priority ($m = 0.80$)
- Balance mix ($m = 0.50$)
- Expansion priority ($m = 0.20$)

Table 4 reports discounted welfare, the main welfare components, and cumulative deforestation units delivered by the model's environmental accounting block.

TABLE 4. SCENARIO OUTCOMES (CENTRAL CALIBRATION, 12-YEAR HORIZON AND AN ANNUAL DISCOUNT RATE OF 8.5%¹⁴)

¹⁴ The 8.5% discount rate corresponds to the Brazilian social discount rate (Nota Técnica SEI n. 19911/2020/ME). We note that this rate is lower than what is traditionally used by, for example, the World Bank, which assumes an annual discount rate of 12% (The World Bank 2005). A sensibility test shows that using 12% instead of 8.5% does not move m^* or flip the scenario ranking. It mainly lowers reported discounted welfare and makes the curve flatter near the top ($W(0.95) - W(1.0) \approx 0.51$ at 12% vs ≈ 0.65 at 8.5%).



Scenario	m	Discounted welfare	Discounted benefits	Discounted fiscal costs	Discounted environmental costs	Total deforestation
Welfare optimal	0.95	13,640	15,744	50	2,054	1,455
Maintenance priority	0.80	13,638	15,744	50	2,057	1,458
Baseline mix	0.50	13,627	15,735	50	2,059	1,460
Expansion priority	0.20	13,604	15,713	50	2,059	1,460

Two patterns stand out. First, welfare ranking is monotone in the direction of maintenance under these scenarios: the optimum lies at the top, followed by the maintenance-heavy rule, the 50-50 mix, and finally the expansion-heavy rule. The welfare loss from *Expansion priority* relative to *Welfare optimal* is on the order of 34 units. The loss from moving from 0.95 to 0.80 is much smaller (2 units), reinforcing that much of the welfare ranking is about avoiding low-maintenance tail rules, not about fine-tuning among already maintenance-dominant allocations.

Second, environmental costs and deforestation move coherently with m in the expected direction: more expansion weight is associated with slightly higher discounted environmental costs and higher cumulative deforestation totals. Comparing the two extreme “priority” rules, *Expansion priority* implies about 2 additional units of cumulative deforestation relative to *Maintenance priority* in the model’s forest-loss accounting (approximately 1,460 vs 1,458). This gap should be read as a model-consistent projection conditional on the reduced-form land-use block, not as a calibrated causal impact evaluation.

It is also useful to note that discounted costs are nearly identical across the four rows in Table 4 (i.e., 50). This pattern reflects the construction of these scenarios as different splits of a common annual envelope, rather than experiments that expand the overall public budget. Differences in welfare therefore come predominantly from benefit



realization under absorption constraints and from environmental externalities, not from a story in which “maintenance is cheaper in fiscal terms” in this accounting.

7. SENSITIVITY AND ROBUSTNESS

The Results section characterizes the welfare–maintenance trade-off under the central calibration. We now test whether the optimal maintenance share m^* and the ordering of the stylized scenarios are robust to alternative discounting and to variation in the quality–access complementarity parameter β_{qa} .

We repeat the central exercise with a 12% annual discount rate, a convention widely used in transport project appraisal and multilateral bank evaluation of road investments (The World Bank 2005), where the rate is often interpreted as an opportunity cost of capital rather than a social time preference rate.

Under a 12% discount rate, the simulation path is unchanged: the same maintenance share m implies the same year-by-year quality, access, deforestation, and undiscounted fiscal and environmental flows. Only the present-value weights applied to those flows change. Discounted welfare therefore falls at every m , for example, from 13,640 at $m^* = 0.95$ under 8.5% to 12,111 at the same allocation under 12%, but the relative ranking of candidate shares is preserved. The welfare-maximizing share remains $m^* = 0.95$, the scenario ordering (welfare-optimal > maintenance-priority > baseline mix > expansion-priority) is unchanged, and cumulative deforestation over the horizon is identical because it is reported as an undiscounted aggregate.

This pattern is expected when a higher discount rate scales down future benefits and costs in a broadly similar way across maintenance-heavy and expansion-heavy rules. The policy-relevant object in this model is the comparative ranking of m , not the absolute level of discounted welfare, which is sensitive to β but not informative about the maintenance–expansion split once ranking is stable. The 12% exercise therefore serves as a conservatism check toward financial-appraisal conventions: even when future flows are discounted more aggressively, the case for a maintenance-tilted interior optimum does not reverse.



We next ask whether m^* remains interior and maintenance-tilted when the strength of quality-access complementarity varies. In the benefit function, access enters as $b_a \cdot a^{\gamma_a} \cdot q^{\beta_{qa}}$: the parameter β_{qa} governs how much realized access benefits depend on maintained road condition. When $\beta_{qa} = 0$, access benefits are effectively independent of quality (additive specification); when β_{qa} approaches 1, access and condition are fully coupled, so connectivity gains from new construction deliver little welfare if the network is allowed to deteriorate.

Holding all other calibration blocks fixed, we sweep $\beta_{qa} \in [0, 1]$ and re-optimize m on the same grid used in the central run. For all $\beta_{qa} < 1$, the optimum lies in the interval $m^* \in [0.76, 0.98]$; $m^* = 1$ (pure maintenance) appears only at the knife-edge case $\beta_{qa} = 1$. At the central calibration $\beta_{qa} = 0.5$, $m^* = 0.95$, matching the baseline reported in Results.

The direction of movement is economically interpretable. Low complementarity (small β_{qa}) weakens the penalty for allocating budget to expansion when condition is poor, because access benefits are less tied to upkeep; the optimum shifts toward $m^* \approx 0.76$, i.e. still maintenance-majority but allowing a larger expansion share than in the central case. High complementarity (large β_{qa}) strengthens the value of preserving condition and makes expansion without maintenance less attractive at the margin; m^* moves toward 0.98, close to—but not at—full maintenance. At no point in this sweep does the optimum flip toward an expansion-dominant corner ($m^* < 0.5$).

Read together with the discount-rate check, these exercises support a narrow but important robustness claim: under the calibrated Amazon panel and literature-anchored parameter blocks, the maintenance–expansion decision is not fragile to moderate changes in discounting or complementarity. What varies is how far from the optimum sits above 50% maintenance, between roughly 76% and 98%, not whether maintenance should dominate the budget envelope.

We also evaluate a 9×8 grid – $b_a \in \{0.50, 0.75, 1.00, 1.25, 1.50, 1.75, 2.00, 2.50, 3.00\}$ and $\beta_{qa} \in \{0.00, 0.10, 0.20, 0.30, 0.40, 0.50, 0.70, 1.00\}$ – with all other parameters at central values and hard caps active. We found that m^* falls as b_a rises (expansion is



more valuable in welfare terms) and rises as β_{qa} rises (maintained quality is more important for capturing access benefits). The central cell ($b_a = 1.50, \beta_{qa} = 0.50$) gives $m^* = 0.95$. Across the rectangle most relevant for appraisal – $b_a \in [1.25, 2.00]$ and $\beta_{qa} \in [0.30, 0.50]$ – m^* spans roughly 0.84-0.96, all interior optima. Corners appear only at the edges of the grid: $m^* = 1$ when b_a is very low (≤ 0.75 with moderate-to-high β_{qa}) or when $\beta_{qa} = 1$ for moderate b_a ; $m^* = 0$ (pure expansion) never appears. Welfare gains of m^* over the pure-expansion benchmark ($m = 0$) range from about 40 to 105 discounted welfare units on the grid, rising with both b_a and β_{qa} because the opportunity cost of wrong-sided allocation is larger when access and complementarity are stronger.

Additionally, holding $\beta_{qa} = 0.5$ and all other central values fixed, we sweep b_a from 1 to 2. The optimum moves smoothly from a maintenance corner at low access weight to a more expansion-tilted interior as b_a rises (Table 5).

TABLE 5. SENSITIVITY ANALYSIS: ACCESS-BENEFIT WEIGHT

b_a	m^*
1-1.1	1 (corner)
1.2	0.99
1.25	0.98
1.4	0.96
1.5	0.95
1.6	0.94
1.8	0.92
2.0	0.90

Table 5 shows that at $b_a \leq 1.10$ the model still prefers 100% maintenance; the interior solution emerges once arterial access is weighted sufficiently above pavement quality. The central choice $b_a = 1.50$ lies in the flat middle of this transition—not at a knife-edge. The welfare gap between m^* and the expansion-priority scenario ($m = 0.2$) stays near 36 discounted units across the sweep, so the ranking of maintenance-first policies is stable even as the exact split shifts.



Finally, while holding deforestation flows and all other parameters at central values, we scale the total damage weight ($SCC_0 + c_{bio}$) from 0 to 3, including balanced splits ($SCC^0 = c_{bio} = total/2$), a carbon-only variant ($c_{bio} = 0$), and threshold scans at totals 1.2-1.8. Total physical deforestation is virtually unchanged across variants (~1,455-1,457 over the 12-year horizon); only the valuation of that flow moves (Table 6).

TABLE 6. SENSITIVITY ANALYSIS: DEFORESTATION PARAMETERS

Scenario	Total damage weight	m^*	Welfare at m^*
Zero damage	0	0.86	15,695
Low/carbon-only	1	0.90-0.91	14,630-14,670
Threshold scan	1.2-1.8	0.91-0.94	13,845-14,462
Central (anchored)	2	0.95	13,640
High damage	3	1	12,613

Based on the results presented in Table 6, we have that the *Higher damage* scenario weights shift m^* toward maintenance (from 86% at zero damage to 100% at the highest scale). The central weight 2.00 sits in the interior of the plausible range: lowering total damage to 1.0 (a carbon-only or high-discount-policy reading) reduces m^* to about 90-91%, while raising it to 3.0 pushes m^* to the maintenance corner (100%). Carbon-only and balanced low-total variants give nearly the same m^* , confirming that composition matters less than the total scale in this specification.

8. LIMITATIONS

As described, the model presented in this study is a simple and computationally feasible model. As a result, there are some important limitations that one must take into account when interpreting the results. Below we highlight the main four limitations:

Partial equilibrium (and no spillovers). The model does not recalculate prices, wages, migration, or land markets inside the simulation. Those forces are not missing entirely:



they enter indirectly through numbers taken from other studies (e.g., elasticities). This limitation is important because although the model can still be useful for comparing maintenance-heavy vs expansion-heavy rules, it is not a full picture of how the whole Amazon economy would adjust. Big economy-wide effects may change how large the gains or losses look. They may also shift the best maintenance share, especially if new roads benefit more from general equilibrium than maintenance (or the opposite), but the direction is not automatic.

Traffic load is treated as fixed. In the model, each corridor's traffic index does not change over time when road quality improves or when new capacity is added. In reality, better or new roads can attract more traffic, and more traffic can wear the road faster. In this sense, the model may miss extra deterioration that would happen if traffic grew. That can make quality paths less realistic and can skew the comparison between spending on maintenance and spending on expansion. It does not mean the results are meaningless, but it means fine details over time should be read cautiously.

Aggregate margin only. The model answers: what share of the budget should go to maintaining existing roads vs new construction, for the current list of federal corridors in the dataset. It does not choose which new roads to build, which to close, or in what order to build projects. The recommendation here is about the balance between two types of spending, not a ranking of individual projects or a network redesign.

No causal attribution. Deforestation in the model follows simple response rules to road variables (a reduced form). The historical clearing measure near each corridor is a summary statistic from satellite data, not proof that each cleared pixel was caused by that specific road. The environmental numbers are consistent stories inside the model, given the parameters we chose. They are not causal impact estimates from satellite data alone. If the true land-use response is very different, environmental totals and the relative cost of expansion could change, which could move the optimal shares somewhat.

Thus, due to the limitations highlighted here, the results should be interpreted jointly with the model scope discussed previously. Because the exercise is partial equilibrium and holds traffic intensity fixed, the quantitative magnitudes are not forecasts of



economy-wide Amazon outcomes under a national infrastructure program. Because the decision margin is aggregate m , the results characterize the split of a fixed budget for a given corridor stock, not which corridor investments should enter a multi-year portfolio. Because deforestation is represented through reduced-form responses, environmental totals should be treated as internally coherent projections under calibrated elasticities, not as causal impact estimates from clearing pixels near roads alone.

Within those boundaries, the central substantive result is stable in direction: a maintenance-heavy interior allocation dominates expansion-heavy rules.

9. CONCLUSION

This study identifies an interior optimum at $m^* = 0.95$, with the full policy ranking detailed in Table 4 demonstrating that the *Welfare optimal* allocation dominates the *Maintenance priority*, *Baseline mix*, and *Expansion priority* scenarios. Based on these findings, we organize the policy discussion around four key dimensions: (i) the budget allocation rule; (ii) specific contexts where expansion remains warranted; (iii) governance and implementation challenges; and (iv) the alignment of these conclusions with broader literature, given model limitations.

1. Budget allocation rule

The central outcome from this study should be interpreted as a fiscal rule posture rather than as a permanent statutory percentage. The implied rule is maintenance-first programming of the federal productive budget: roughly nine-to-one maintenance versus new construction in the main calibration, reflecting the optimum $m^* = 0.95$. This is not an argument for $m = 1$: the model's own welfare grid shows $m = 1$ is slightly inferior to $m^* = 0.95$, because a small expansion margin still purchases access benefits that are not fully substitutable by maintenance of existing links alone, subject to the model's spatial representation.

In practical budget terms, the implication is closer to treating routine and restorative maintenance as the core recurrent investment of federally administered highways in



the Legal Amazon, and treating lane-kilometer expansion as a residual margin that must clear engineering absorption and environmental risk tests. This aligns with Table 4's lesson that most of the welfare loss relative to *Welfare optimal* is concentrated when the budget tilts toward *Expansion priority* ($m = 0.20$), while differences among already maintenance-heavy rules are smaller (e.g., *Welfare optimal* vs *Maintenance priority*).

Finally, the model's hard caps imply a cap-aware allocation norm. In our welfare accounting, spending beyond corridor physical or institutional capacity is socially unproductive yet it remains fiscally costly. Translated into governance, budget authorizations must be tied directly to realistic, executable engineering programs. Planners must avoid the common bureaucratic pitfall of automatically reallocating "unspent maintenance funds" into expansion projects when those projects face binding institutional ceilings or high environmental risks.

2. Where expansion may still be justified

Even under $m^* = 0.95$, expansion is not zero. Expansion remains relatively attractive where access returns are strong and where new construction can be absorbed without a large share of spending wasted against tight new-road caps—in particular where protected-area overlap does not mechanically compress the permissible expansion envelope in the modelling framework.

Three policy-readable conditions summarize the model's logic (without duplicating project-selection optimization, which is outside the model's choice set):

- Meaningful access payoffs consistent with a network-improvement interpretation of federal corridors (as opposed to settings where rural feeder expansion delivers negligible productivity returns).
- Credible physical execution of new-capacity works (so marginal public spending translates into productive connectivity rather than deadweight against caps).
- Institutional and environmental safeguards aligned with lower marginal forest-risk segments, recognizing that the overlap rule in the database is federal conservation units geometry only and must be complemented in real licensing by broader conservation realities.



This bridges directly to Table 4's co-movement: *Expansion priority* is associated with higher modelled cumulative deforestation than *Maintenance priority*. The policy implication is not "never expand," but "expand selectively," and treat large expansion packages as high-risk fiscal-environmental portfolios unless mitigation and enforcement are commensurate.

3. Governance and implementation implications

Budget governance. Maintenance-dominance requires stable multi-year programming for rehabilitation cycles, not treating maintenance as the residual margin when discretionary announcements expand. That institutional point is conceptually separate from the numerical $m^* = 0.95$, but it is what makes a maintenance-first implementable in real public finance.

Monitoring and data updating. Because m^* lies in a high-maintenance region under β_{qa} robustness (typically [0.76, 0.98] for $\beta_{qa} < 1$), coarse robustness to the broad tilt is plausible, while fine calibration can still move within that band. Operationally, this argues for annual refresh of traffic proxies, condition metrics (ideally beyond the current status-based proxy), and deforestation/enforcement context—not because the model claims year-by-year forecasting precision, but because the planner problem is data-sensitive whenever political debate focuses on shifting m within the high-maintenance range.

Environmental governance realism. The enforcement variable used to build effective protection is a coarse federal proxy; implementation cannot equate IBAMA intensity with total environmental capacity. Effective governance for expansion combines remote detection, field enforcement, sanctions that bind, and state agency complements—especially in states where federal counts understate total enforcement effort.

Intergovernmental coordination. The model joins federal corridors to state controls; real programs still require licensing, environmental compensation, and road safety institutions beyond a federal envelope split exercise.

4. Alignment broader literature (in light of model limitations)



The model is consistent with transport asset-management logic that service quality on an existing network is a stock requiring recurrent investment. Table 4's deterioration of outcomes toward *Expansion priority* parallels the policy risk of under-maintaining a paved network while prioritizing new kilometers.

The model is also consistent with Amazon evidence emphasizing that frontier connectivity can carry large environmental externalities, in the sense that the model does not produce expansion-first optima under the calibrated welfare and damage structure.

Finally, the model is consistent with literatures where monitoring and enforcement matter for whether conservation regimes bind, at least to the limited extent that our model includes measurable governance covariates and protection overlap penalties on expansion ceilings.

All in all, the results found in this study show that the policy implications are consistent with the model's maintenance-expansion margin and the ranking in Table 4 around $m^* = 0.95$.



REFERENCES

- Allen, T., & Arkolakis, C. (2014). Trade and the Topography of the Spatial Economy. *The Quarterly Journal of Economics*, 129(3), 1085–1140.
- Araujo, R., Assunção, J., & Bragança, A. A. (2023). The Effects of Transportation Infrastructure on Deforestation in the Amazon: A General Equilibrium Approach. The World Bank. <https://doi.org/10.1596/1813-9450-10415>
- Araujo, R., Assunção, J., & Bragança, A. (2025). Transportation infrastructure and deforestation in the Amazon. *Journal of Development Economics*, 177, 103559. <https://doi.org/10.1016/j.jdeveco.2025.103559>
- Araujo, R., Bragança, A., & Assunção, J. Accessibility in the Legal Amazon: Measuring Market Access. *Amazônia 2030*, 2022.
- Asher, S., Garg, T., & Novosad, P. (2020). The Ecological Impact of Transportation Infrastructure. *The Economic Journal*, 130(629), 1173–1199. <https://doi.org/10.1093/ej/ueaa013>
- Assunção, J., Gandour, C., & Rocha, R. (2023). DETER-ing Deforestation in the Amazon: Environmental Monitoring and Law Enforcement. *American Economic Journal: Applied Economics*, 15(2), 125–156. <https://doi.org/10.1257/app.20200196>
- Barber, C. P., Cochrane, M. A., Souza, C. M., & Laurance, W. F. (2014). Roads, deforestation, and the mitigating effect of protected areas in the Amazon. *Biological Conservation*, 177, 203–209. <https://doi.org/10.1016/j.biocon.2014.07.004>



Botelho, J., Costa, S. C. P., Ribeiro, J. G., & Souza, C. M. (2022). Mapping Roads in the Brazilian Amazon with Artificial Intelligence and Sentinel-2. *Remote Sensing*, 14(15), 3625. <https://doi.org/10.3390/rs14153625>

Cezne, E., & Otsuki, K. (2025). The Amazon's road to growth? Infrastructural imaginaries of Brazil's BR-319 Highway. *Economy and Society*, 54(4), 620–646.

<https://doi.org/10.1080/03085147.2025.2588934>

DNIT (Departamento Nacional de Infraestrutura de Transportes). (2022). Volume médio diário anual (VMDa), modelagem 2022 [Database / geographic layer]. DNIT. Retrieved April 15, 2026.

DNIT (Departamento Nacional de Infraestrutura de Transportes). (2025a). Boletim Infraestrutura. Centro de Estudos em Infraestrutura FGV IBRE. Junho 2025. Available from: https://www.gov.br/dnit/pt-br/assuntos/planejamento-e-pesquisa/custos-referenciais/boletim-da-infraestrutura/publicacoes-1/2025-06_boletim-infraestrutura_edicao-02.pdf

DNIT (Departamento Nacional de Infraestrutura de Transportes). (2025b). Sistema Nacional de Viação (SNV): Federal road network [Spatial data, snv_202507a extract]. Retrieved April 15, 2026.

DNIT. (Departamento Nacional de Infraestrutura de Transportes). (2026). Relatório de Gestão 2025. Ministério dos Transportes, Brasília, Brasil. Available from:

<https://www.gov.br/dnit/pt-br/acesso-a-informacao/auditoria/relatorio-de-gestao>.



Donaldson, D. (2018). Railroads of the Raj: Estimating the Impact of Transportation Infrastructure. *American Economic Review*, 108(4-5), 899-934.

<https://doi.org/10.1257/aer.20101199>

Fearnside, P.M. 1987. Causes of Deforestation in the Brazilian Amazon. pp. 37-61 In: R.F. Dickinson (ed.) *The Geophisiology of Amazonia: Vegetation and Climate Interactions*. John Wiley & Sons, New York. 526 pp.

Guayasamin JM, Ribas CC, Carnaval AC, Carrillo JD, Hoorn C, Lohmann LG, Riff D, Ulloa Ulloa C, Albert JS. 2021. Chapter 2: Evolution of Amazonian Biodiversity. In: Nobre C, Encalada A, Anderson E, Roca Alcazar FH, Bustamante M, Mena C, Peña-Claros M, Poveda G, Rodriguez JP, Saleska S, Trumbore S, Val AL, Villa Nova L, Abramovay R, Alencar A, Rodríguez Alza C, Armenteras D, Artaxo P, Athayde S, Barretto Filho HT, Barlow J, Berenguer E, Bortolotto F, Costa FA, Costa MH, Cui N, Fearnside PM, Ferreira J, Flores BM, Frieri S, Gatti LV, Guayasamin JM, Hecht S, Hirota M, Hoorn C, Josse C, Lapola DM, Larrea C, Larrea-Alcazar DM, Lehm Ardaya Z, Malhi Y, Marengo JA, Melack J, Moraes R M, Moutinho P, Murmis MR, Neves EG, Paez B, Painter L, Ramos A, Rosero-Peña MC, Schmink M, Sist P, ter Steege H, Val P, van der Voort H, Varese M, Zapata-Ríos G (Eds). *Amazon Assessment Report 2021*. United Nations Sustainable Development Solutions Network, New York, USA. Available from <https://www.theamazonwewant.org/spa-reports/>. DOI: 10.55161/CZWN4679

Hanusch, M. (Ed.). (2023). *A Balancing Act for Brazil's Amazonian States: An Economic Memorandum*. The World Bank. <https://doi.org/10.1596/978-1-4648-1909-4>

Herrera Dappe, M.; Lebrand, M.; Stokenberga, A. (2024). *Shrinking Economic Distance: Understanding How Markets and Places Can Lower Transport Costs in*



Developing Countries. Sustainable Infrastructure Series. World Bank.

<http://hdl.handle.net/10986/42061> License: CC BY 3.0 IGO.

IBAMA (Instituto Brasileiro do Meio Ambiente e dos Recursos Naturais Renováveis). (2023). Environmental infraction notices (Autuações ambientais) [Administrative records]. IBAMA. Retrieved May 1, 2026.

IBGE (Instituto Brasileiro de Geografia e Estatística). (2023). Regional accounts of Brazil: Gross value added by state [Data set]. IBGE. Retrieved May 1, 2026.

IBGE (Instituto Brasileiro de Geografia e Estatística). (2024). Amazônia Legal (Legal Amazon boundaries) [Spatial data]. IBGE. Retrieved April 20, 2026

IBGE (Instituto Brasileiro de Geografia e Estatística). (2025). Estimated resident population by municipality, 2025 [Data set, SIDRA Table 6579]. IBGE SIDRA. Retrieved April 15, 2026, from <https://sidra.ibge.gov.br/tabela/6579>.

ICMBio (Instituto Chico Mendes de Conservação da Biodiversidade). (2026). Federal conservation units – geospatial boundaries [Spatial data]. Retrieved April 15, 2026

Lanzer, B. N. (2011). Ensaio sobre o custo marginal do financiamento público no Brasil. <https://lume.ufrgs.br/handle/10183/49928>

Lapola, D. M., Pinho, P., Barlow, J., Aragão, L. E. O. C., Berenguer, E., Carmenta, R., Liddy, H. M., Seixas, H., Silva, C. V. J., Silva-Junior, C. H. L., Alencar, A. A. C., Anderson, L. O., Armenteras, D., Brovkin, V., Calders, K., Chambers, J., Chini, L., Costa, M. H., Faria, B. L., ... Walker, W. S. (2023). The drivers and impacts of Amazon forest degradation. *Science*, 379(6630), eabp8622. <https://doi.org/10.1126/science.abp8622>



Laurance, W. F., Goosem, M., & Laurance, S. G. W. (2009). Impacts of roads and linear clearings on tropical forests. *Trends in Ecology & Evolution*, 24(12), 659–669.

<https://doi.org/10.1016/j.tree.2009.06.009>

Malhi Y, Melack J, Gatti LV, Ometto J, Kesselmeier J, Wolff S, Aragão LEO, Costa M, Saleska S, Pangala SR, Basso LS, Rizzo L, Araùjo AC, Restrepo-Coupe N, Junior CHLS. 2021. Chapter 6: Biogeochemical Cycles of the Amazon. In: Nobre C, Encalada A, Anderson E, Roca Alcazar FH, Bustamante M, Mena C, Peña-Claros M, Poveda G, Rodriguez JP, Saleska S, Trumbore S, Val AL, Villa Nova L, Abramovay R, Alencar A, Rodríguez Alza C, Armenteras D, Artaxo P, Athayde S, Barretto Filho HT, Barlow J, Berenguer E, Bortolotto F, Costa FA, Costa MH, Cuvi N, Fearnside PM, Ferreira J, Flores BM, Frieri S, Gatti LV, Guayasamin JM, Hecht S, Hirota M, Hoorn C, Josse C, Lapola DM, Larrea C, Larrea-Alcazar DM, Lehm Ardaya Z, Malhi Y, Marengo JA, Melack J, Moraes R M, Moutinho P, Murmis MR, Neves EG, Paez B, Painter L, Ramos A, Rosero-Peña MC, Schmink M, Sist P, ter Steege H, Val P, van der Voort H, Varese M, Zapata-Ríos G (Eds). *Amazon Assessment Report 2021*. United Nations Sustainable Development Solutions Network, New York, USA. Available from

<https://www.theamazonwewant.org/spa-reports/>. DOI: 10.55161/TAKR3454

MapBiomas. (2026). Annual deforestation layers, Brazil, 2014-2024 [Geospatial data].

MapBiomas Project. Retrieved April 15, 2026.

Matricardi, E. A. T., Skole, D. L., Costa, O. B., Pedlowski, M. A., Samek, J. H., & Miguel, E. P. (2020). Long-term forest degradation surpasses deforestation in the Brazilian Amazon. *Science*, 369(6509), 1378–1382. <https://doi.org/10.1126/science.abb3021>



Pellegrina, H. S., & Sotelo, S. (2025). Migration, Specialization, and Trade: Evidence from Brazil's March to the West. *Journal of Political Economy*, 133(12), 3993-4049. <https://doi.org/10.1086/738147>

Pfaff, A. S. P. (1999). What Drives Deforestation in the Brazilian Amazon?: Evidence from Satellite and Socioeconomic Data. *Journal of Environmental Economics and Management*, 37(1), 26-43. <https://doi.org/10.1006/jeem.1998.1056>

Programa das Nações Unidas para o Desenvolvimento (PNUD), Instituto de Pesquisa Econômica Aplicada (IPEA), & Fundação João Pinheiro (FJP). (2022). Atlas do Desenvolvimento Humano no Brasil [Data set]. PNUD Brasil. <https://www.atlasbrasil.org.br/>

Skole, D., & Tucker, C. (1993). Tropical Deforestation and Habitat Fragmentation in the Amazon: Satellite Data from 1978 to 1988. *Science*, 260(5116), 1905-1910. <https://doi.org/10.1126/science.260.5116.1905>

Soares-Filho, B. S., Nepstad, D. C., Curran, L. M., Cerqueira, G. C., Garcia, R. A., Ramos, C. A., Voll, E., McDonald, A., Lefebvre, P., & Schlesinger, P. (2006). Modelling conservation in the Amazon basin. *Nature*, 440(7083), 520-523. <https://doi.org/10.1038/nature04389>

Souza, L. H. C. D., & Silva, S. D. (2022). The Economic and Social Impacts of Repaving BR 319 and their Respective Environmental Developments. *International Journal of Advanced Engineering Research and Science*, 9(12), 203-219. <https://doi.org/10.22161/ijaers.912.23>



The World Bank. (2005). A Framework for the Economic Evaluation of Transport Projects. Notes on the Economic Evaluation of Transport Projects. Transport Note N. TRN-5. Washington, DC.

Vilela, T., Malky Harb, A., Bruner, A., Laísa da Silva Arruda, V., Ribeiro, V., Auxiliadora Costa Alencar, A., Julissa Escobedo Grandez, A., Rojas, A., Laina, A., & Botero, R. (2020). A better Amazon road network for people and the environment. *Proceedings of the National Academy of Sciences*, 117(13), 7095–7102.

Wenzel, F. (2026, April 20). Brazil taps legal loophole to issue bids for Amazon ‘tipping point’ road. *Mongabay*. <https://news.mongabay.com/2026/04/brazil-taps-legal-loophole-to-issue-bids-for-amazon-tipping-point-road/>

Zhou, D., Xiao, J., Liu, S., Huang, L., Zhang, L., Fan, J., & Zhao, S. (2026). Global impacts of transportation infrastructure on forest degradation and loss. *Nature Communications*, 17(1), 2339. <https://doi.org/10.1038/s41467-026-69150-4>