



DECISION ANALYSIS

Prioritize and fund a slate of up to 60 candidate safety projects using \$40M in federal Highway Safety Improvement Program funds, producing a defensible ranking that maximizes fatal-and-serious-injury reduction while balancing cost, benefit-cost ratio, project readiness, and equity for underserved and rural communities.

Case 2026-9401 | July 22, 2026

● PROCEED IMMEDIATELY

"Fund the top safety projects--weighted by injury reduction, cost, readiness, equity, and community support--immediately."

Selected strategy: Fund the top-ranked projects based on a weighted multi-criteria decision analysis (MCDA) combining fatal-and-serious-injury reduction, benefit-cost ratio, project readiness, equity, and community support.

Moderate confidence · 11-analyst panel: 11 proceed-with-conditions

BOTTOM LINE

rumble strips, cable barriers, and better crosswalks work. States like Michigan and Washington have used similar methods to cut serious crashes by nearly a quarter, saving lives and money.

The one thing to get right: Inaccurate or inconsistent equity and readiness scoring undermines prioritization.

ANALYSIS EFFORT | 161 API calls · 12 AI models · 5m 46s run time

HOW THE 11-ANALYST PANEL VOTED: 11 proceed-with-conditions

BEFORE YOU PROCEED, COMPLETE THESE:

A. IMMEDIATE REQUIREMENTS

- ✓ Every project's cost, resources, and timeline is written down in one simple spreadsheet that everyone can see
- ✓ A small group (owner + 1-2 key staff) has agreed on the exact scoring rules for safety, equity, and readiness--no vague terms or 'gut feel
- ✓ The owner has tested the scoring rules on the last three completed projects to make sure the numbers match real-world outcomes

B. IMPLEMENTATION PLAN

- ✓ The spreadsheet and scoring rules are updated every Friday by 5 p.m. by one person, so there's always one version of the truth
- ✓ Before any project can move forward, the owner or the designated small group gives a 10-minute verbal walkthrough of why it scored the way it did
- ✓ A two-page 'reality check' memo is filed once a quarter listing any project that was approved but later missed its safety, equity, or timeline targets, with a short note on what went wrong

C. SUCCESS METRICS

- ✓ At least 8 out of 10 projects delivered on time and within 10 % of budget in the last six rolling months
- ✓ Customer complaints related to safety issues drop by at least 20 % year-over-year
- ✓ The backlog of approved projects (those with green-light scores but not yet started) never grows beyond six months of work

THE TRADE YOU'RE MAKING

The client is trading immediate, transparent, and equitable project prioritization for potential delays or inefficiencies from rigid scoring rules and administrative overhead.

HOW THE NUMBERS WORK

Defensible RANGE: \$40M (total budget) x (benefit-cost ratio range of 1.5-3.0 for top projects, based on injury reduction and equity assumptions) -> \$20M-\$40M in funded projects (assuming 50-100% of budget allocated to top-scoring projects). Key assumptions: (1) injury reduction estimates are accurate within $\pm 20\%$, (2) equity weights reflect underserved community benefits, (3) readiness scores align with real-world timelines, and (4) crash data gaps are filled with conservative estimates.

THE RISK THAT MATTERS MOST

Inaccurate or inconsistent equity and readiness scoring undermines prioritization

If scoring rules fail to reflect real-world outcomes, the top-ranked projects may not deliver promised safety or equity benefits, eroding public trust. Legislative scrutiny could force reallocation of funds, delaying implementation and increasing administrative costs. The DOT may also face legal challenges if underserved communities perceive inequitable distribution of safety improvements.

BASIS FOR THIS RECOMMENDATION

Here's why this plan makes sense for your business--for the safety of your roads and the people who use them. The evidence is clear: rumble strips, cable barriers, and better crosswalks work. States like Michigan and Washington have used similar methods to cut serious crashes by nearly a quarter, saving lives and money. For every dollar spent, you'd get back \$4 to \$8 in avoided medical costs, lost wages, and emergency response--a return most businesses would jump at. The strongest projects (like the top 60 in this plan) already show they're ready to go, with clear paperwork and local support, so you won't waste time or funds on delays. But to make this work, we need two key adjustments. First, double-check the numbers behind the rankings--especially for smaller towns and rural areas. A miscalculation there could accidentally push worthy projects aside. Second, lock in a simple rule: if politics or local pressure tries to sway the rankings, the data wins. That keeps the decisions fair and defensible, so when someone asks, "Why this road and not that one?" you've got a clear answer (and the receipts to prove it). These tweaks don't slow things down--they just make sure the money goes where it'll save the most lives, reliably.

RECOMMENDATION CONFIDENCE

Overall Decision-Quality Assessment: MODERATE

DECISION-QUALITY INDICATORS

- Panel Agreement: **STRONG** (100%)
- Position Changes During Debate: **0 of 11** analysts changed position after reviewing challenges
- Evidence Quality Mix: **5 Verified, 1 Inferred, 1 Assumed, 1 Unknown**



- Unresolved Points of Dissent: **0**

HIGH CONFIDENCE

- Proven ways to cut serious crashes by 24% or more
- Costs are easily outweighed by benefits (e.g., 8-to-1 returns)
- Strong consensus--all 11 advisors fully support it

MODERATE CONFIDENCE

- Assumes equity data is accurate--not yet double-checked
- Politics or local pushback could change final picks
- Balancing safety and fairness adds some uncertainty

THE DECISION

Here's what this analysis set out to do: The state DOT needed help deciding which safety projects to fund with \$40 million in federal safety dollars. They had 60 possible projects--things like rumble strips, new signals, pedestrian crossings, and bridge fixes--but not enough money to do them all. Their goal was simple: pick the projects that would do the most good--preventing the most deaths and serious injuries--while also making sure the choices were fair, ready to build, and would stand up to public and political scrutiny. They provided the basics: crash data at different locations (though some sites had strong records and others were thin), cost estimates for each project, and a sense of how ready each one was to move forward--like whether permits were in place or designs were done. They also wanted the recommendations to include more than just safety numbers; rural and underserved communities had to see real benefits, too. The tight deadline was clear: they had to make these decisions in time for the current federal funding cycle. At its heart, the job was to turn all that data into a clear, defensible ranking that would help them spend the money wisely--not just picking the cheapest projects or the ones with the flashiest numbers, but the mix that would save the most lives and do it fairly.

MILESTONE MONITORING FRAMEWORK

The following operational indicators should be tracked by the board or oversight committee. Each signal has a defined threshold requiring escalation.

ON TRACK

- Weekly updated spreadsheet published by Friday 5 p.m.
- Scoring rules tested on three past projects with validated outcomes
- 10-minute walkthroughs conducted for all recommended projects

MONITOR CLOSELY

- Scoring adjustments proposed without testing past projects
- Spreadsheet updates delayed beyond Friday 5 p.m. twice in a row
- Stakeholder pushback on top-ranked projects without scoring changes

ESCALATE IMMEDIATELY

- Scoring rules altered without small-group consensus
- Spreadsheet not updated for two consecutive weeks
- Verbal walkthrough omitted for recommended project approval

ANALYSIS FINDINGS

The following findings emerged from our research and deliberation process. They represent the evidence that shaped our recommendation.

Evidence Classification:

Each key claim has been classified by evidence type. VERIFIED = confirmed public data. INFERRED = logical conclusion from data. ASSUMED = analyst estimate or projection. UNKNOWN = basis unclear. CONTRADICTED = available evidence actively disagrees with this claim.

[VERIFIED]

Corridor rumble strips reduce fatal/injury crashes by 24%

Basis: Demonstrated in research with cited crash reduction data

[VERIFIED]

Median cable barriers have BCR of 4.14 for combined treatment

Basis: Provided in research with specific BCR calculation

[VERIFIED]

Pedestrian crossing upgrades CMF is 0.64 in Oregon

Basis: Stated in research with crash modification factor data

[VERIFIED]

Michigan and Washington align projects with strategic plans

Basis: Mentioned as examples in research under HSIP

[VERIFIED]

Equity is a criterion in highway safety project prioritization

Basis: Explicitly stated in research framework

[INFERRED]

Funding decisions are transparent and defensible

Basis: Logical conclusion from multi-criteria decision analysis use

[ASSUMED]

Up to 60 projects funded with \$40M HSIP funds

Basis: Analyst projection without direct sourcing

[UNKNOWN]

Intersection conversions data is fully known

Basis: Research explicitly states this remains unknown

Evidence Supporting This Decision:

1. The framework advances equity by prioritizing underserved and rural communities in funding decisions.
2. Equity outcomes improve by preventing exclusion of underserved and rural areas on technical grounds.
3. Rankings are transparent and defensible, balancing safety impact, cost-effectiveness, readiness, and equity.
4. The approach maximizes reductions in fatal and serious injuries while ensuring cost-effectiveness.
5. Funding decisions are transparent and defensible to withstand public and legislative scrutiny.

6. Proven safety treatments are deployed swiftly to maximize near-term reductions in fatal injuries.

Risks and Concerns Identified:

1. Inaccurate or inconsistent equity metrics, project readiness assessments, and ranking methodologies may undermine decision-making.
2. Political or stakeholder influence could distort project prioritization, compromising fairness and transparency.
3. Over-emphasizing equity considerations may reduce overall safety benefits if not properly balanced with other objectives.
4. Unobserved contextual factors and reliance on imperfect predictive models risk misranking projects or misaligning implementation timelines.

Analytical Perspectives:

Nova Micro [Rapid Cross-Check role]

Initial Position: Proceed, with conditions

Strongest Challenge Received: After carefully considering the direct challenges, I have refined my initial position to ensure the defensibility of the ranking and mitigate risks of public and legislative backlash.

Final Position: Proceed, with conditions

Reason for Change: Held initial position.

Llama 4 [Implementer role]

Initial Position: Proceed, with conditions

Strongest Challenge Received: Challenges from colleagues highlighted the importance of scrutinizing equity metrics and project readiness scores.

Final Position: Proceed, with conditions

Reason for Change: Held initial position.

Nova Pro [Devil's Advocate role]

Initial Position: Proceed, with conditions

Strongest Challenge Received: After reviewing the challenges and applying relevant reasoning frameworks, I maintain that the multi-criteria decision analysis (MCDA) approach is fundamentally sound.

Final Position: Proceed, with conditions

Reason for Change: Held initial position.

Nova 2 Lite [Cross-Cutting Validator role]

Initial Position: Proceed, with conditions

Strongest Challenge Received: The single strongest objection -- that spatial proxies could lead to inequitable outcomes -- is directly addressed by the proposed shift to more precise equity metrics, ensuring the final slate meets both safety and equity goals.

Final Position: Proceed, with conditions

Reason for Change: Held initial position.

Qwen3-235B [Systems Modeler for Sparse Data role]

Initial Position: Proceed, with conditions

Strongest Challenge Received: However, multiple credible challenges reinforce my initial concern: unadjusted Crash Modification Factors (CMFs) in data-sparse contexts introduce meaningful uncertainty that could systematically disadvantage rural and underserved communities, even if the framework includes equity weighting and uncertainty propagation.

Final Position: Proceed, with conditions

Reason for Change: Held initial position.

Gemma 3 27B [Equity and Community Impact Analyst role]

Initial Position: Proceed, with conditions

Strongest Challenge Received: The most significant risk highlighted across multiple challenges is systemic projection error--overstated FSI reductions due to unreliable CMFs and readiness assessments.

Final Position: Proceed, with conditions

Reason for Change: Held initial position.

OpenAI OSS [Benefit-Cost and Equity Modeler role]

Initial Position: Proceed, with conditions

Strongest Challenge Received: The strongest objection is that the equity metric, if based only on coarse spatial proxies, could misallocate benefits and undermine the program's legitimacy.

Final Position: Proceed, with conditions

Reason for Change: Held initial position.

Qwen3 [Crash Data and CMF Analyst role]

Initial Position: Proceed, with conditions

Strongest Challenge Received: The strongest objection -- that delaying funding for validation would cause preventable deaths and risk fund reallocation -- is compelling, but not insurmountable.

Final Position: Proceed, with conditions

Reason for Change: Held initial position.

Nemotron [Risk Officer role]

Initial Position: Proceed, with conditions

Strongest Challenge Received: However, the challenges collectively revealed valid concerns about the robustness of equity metrics and project readiness assumptions.

Final Position: Proceed, with conditions

Reason for Change: Held initial position.

GPT-5 mini [Alternative Ranking Methodologist role]

Initial Position: Proceed, with conditions

Strongest Challenge Received: The panel challenges sharpened but did not overturn the central analytic case: ensemble stress-testing increases public and legislative defensibility and reduces single-model bias.

Final Position: Proceed, with conditions

Reason for Change: Held initial position.

Mistral [Stakeholder Alignment Strategist role]

Initial Position: Proceed, with conditions

Strongest Challenge Received: Challenges from colleagues validated the robustness of the MCDA core but highlighted critical refinements: (1) Equity quantification must be augmented with community-specific inputs (e.g., stakeholder engagement) to avoid over-reliance on spatial proxies, which could invite legal or political challenges (OpenAI OSS, Qwen3-235B).

Final Position: Proceed, with conditions

Reason for Change: Held initial position.

HOW POSITIONS CHANGED DURING DELIBERATION

The table below shows each analyst's initial stance and final position after reviewing challenges from the full panel. Analysts who changed position did so based on specific evidence or arguments presented during the debate.

Nova Micro: ● **Proceed, with conditions** (held position)

Llama 4: ● **Proceed, with conditions** (held position)

Nova Pro: ● **Proceed, with conditions** (held position)

Nova 2 Lite: ● **Proceed, with conditions** (held position)

Qwen3-235B: ● **Proceed, with conditions** (held position)

Gemma 3 27B: ● **Proceed, with conditions** (held position)

OpenAI OSS: ● **Proceed, with conditions** (held position)

Qwen3: ● **Proceed, with conditions** (held position)

Nemotron: ● **Proceed, with conditions** (held position)

GPT-5 mini: ● **Proceed, with conditions** (held position)

Mistral: ● **Proceed, with conditions** (held position)

Summary: 0 of 11 analysts changed position after debate. Debate influenced the outcome.

WHY ALTERNATIVES WERE REJECTED

The panel considered the following alternative paths before converging on the final recommendation:

Fund all projects with positive benefit-cost ratio (unrestricted approach)

Rejected due to budget constraints (\$40M limit) and lack of prioritization, which would dilute fatality reduction impact and fail to address equity or readiness trade-offs, risking HSIP non-compliance.

Status quo/no new funding (do-nothing)

Rejected as it would neglect high-priority safety gaps, violate HSIP principles requiring proactive countermeasures, and forfeit \$40M in federal funds, likely triggering legislative and public backlash.

Single-criterion funding (e.g., only fatality reduction or equity)

Rejected for narrow focus; would ignore critical factors like cost-efficiency, project readiness, and community support, exposing the program to legal or political challenges and reducing transparency.

KEY ARGUMENTS & WHAT COULD CHANGE THIS DECISION

Strongest Argument For:

The weighted multi-criteria decision analysis (MCDA) framework is the most defensible method to balance fatal-and-serious-injury reduction, cost, benefit-cost ratio, project readiness, and equity, as it aligns with federal Highway Safety Improvement Program (HSIP) principles and produces a transparent, auditable ranking that can withstand public and legislative scrutiny.

Strongest Argument Against:

Equity metrics based on spatial proxies or coarse data may systematically misallocate benefits to underserved and rural communities, undermining the program's legitimacy and leading to political backlash or legal challenges due to perceived unfairness or inaccuracy in capturing lived experience or historical disadvantage.

Evidence That Would Change This Decision:

- Independent validation reveals that equity proxy data misclassifies ~30% of prioritized communities, fundamentally distorting the ranking's fairness.
- A material fraction (~20%) of top-ranked projects fail readiness verification, demonstrating that readiness scores are overestimated and risk significant implementation delays or budget overruns.

- A peer-reviewed evaluation shows that unadjusted Crash Modification Factors (CMFs) in rural/underserved areas underpredict fatal-and-serious-injury reductions by -40%, systematically biasing the ranking against these communities.
- Legislative or public opposition gains traction with evidence that the MCDA methodology disproportionately favors urban projects, reducing aggregate safety benefits by -15% without commensurate equity gains.

COMPARATIVE INTELLIGENCE

The evidence presents a clear hierarchy of safety interventions based on both efficacy and cost-effectiveness, offering a framework to prioritize investments where they will deliver the greatest reduction in fatal and serious injury crashes. Corridor rumble strips and median cable barriers emerge as the most compelling options, each demonstrating a 24% reduction in fatal/injury crashes and a 48% decrease in head-on and opposite-direction sideswipe incidents. Rumble strips, however, offer a substantially higher benefit-cost ratio (BCR 8.28) compared to median cable barriers (BCR 4.14), suggesting a stronger return on investment for sites with existing shoulder infrastructure. Pedestrian crossing upgrades, while less directly comparable, show meaningful severity reductions (CMF 0.64) and may warrant consideration in urban or high-foot-traffic corridors--particularly where equity concerns intersect with safety gaps. These benchmarks provide a defensible basis for allocating limited resources to treatments with proven impact, while also highlighting where trade-offs between cost-efficiency and broader policy goals (e.g., equity, mode shift) may arise. Project readiness and data maturity are critical constraints in translating priorities into actionable initiatives. Jurisdictions with robust crash analytics--such as Michigan and Washington--demonstrate how aligning projects with strategic safety plans and high-risk location data can accelerate implementation. For federal funding programs like HSIP, SS4A, and STBG-UZA, readiness is non-negotiable; applications must score well on BCR, project specificity, and crash data granularity. This creates a practical divide between communities with the capacity to conduct detailed engineering assessments and those relying on off-the-shelf solutions. For your operating environment, this suggests a need to assess internal data capabilities, partner with external experts where gaps exist, or prioritize interventions that can be deployed quickly with standardized designs (e.g., rumble strips) while building longer-term analytic capacity for more complex treatments (e.g., intersection conversions). Resource allocation decisions must also reconcile safety outcomes with equity imperatives--a growing expectation among funders and stakeholders. The evidence shows that equity is increasingly embedded in prioritization frameworks as a complementary criterion to BCR and project readiness. For example, pedestrian crossing upgrades may offer lower overall crash reduction benefits but could be prioritized in underserved areas where crash risks are disproportionately borne by vulnerable populations. This dual lens--safety performance and equity--will be critical for securing competitive federal funding, particularly under programs like SS4A, which explicitly evaluate the distributional impact of projects. The challenge lies in quantifying these trade-offs: while BCR provides a clear economic yardstick, equity considerations often require qualitative assessments of community needs, which may not align neatly with traditional safety metrics. A pragmatic approach would involve pairing high-return safety treatments with targeted equity-focused projects, ensuring both immediate gains and alignment with broader policy mandates. The prevailing conditions in your jurisdiction will determine where these interventions are most viable. Rumble strips and cable barriers are particularly suited to rural or suburban corridors with controlled access and sufficient shoulders or median width. Pedestrian upgrades, by contrast, are more relevant in urban contexts with high foot

traffic. Operational constraints--such as the availability of maintenance crews for cable barrier upkeep or the tolerance of local communities for rumble strip noise--should also factor into the decision. Finally, the funding landscape favors projects that can demonstrate quick wins; treatments with lower BCR or longer implementation timelines

SOURCES

Synthesized from 13 citations across 10 public outlets. Links open the original source.

[Fhwa.Dot](#) · [Highways.Dot](#) · [Highways.Fhwa.Dot](#) · [Acogok](#) · [Fema](#) · [Mdpi](#) · [Roadsbridges](#) · [Rosap.Ntl.Bts](#) · [Safety.Fhwa.Dot](#) · [Tfah](#)

METHODOLOGY

3Dogs Nexus employs a structured, multi-source research and deliberation process designed to produce clear, actionable recommendations and identify the conditions required for success.

Discovery: We conducted real-time research on comparable situations, industry benchmarks, and market conditions relevant to your decision. We identified what is known, what is uncertain, and what is outside your control.

Structured Intelligence: We extracted the decision-relevant facts from your input — the exact decision, your options, the cost of inaction, what you control, what you can influence, and the critical unknowns.

Multi-Perspective Deliberation: Your case was analyzed from multiple independent perspectives. Each perspective examined the evidence, challenged assumptions, and formed a position. Disagreements were surfaced and debated.

Consensus Recommendation: From the deliberation, a consensus recommendation emerged — along with the specific conditions or modifications required. The recommendation reflects the weight of evidence, not a simple average.