International Journal of Economics, Commerce and Management
United Kingdom ISSN 2348 0386 Vol. 14, Issue 10, Oct 2026

https://ijecm.co.uk/

DRIVER IMPAIRMENT AS AN UNPRICED EXTERNALITY IN FREIGHT TRANSPORTATION COST MODELS

Lebedev Kirill Olegovich
Professor at the Department of Transport Information Systems
Academy of Advanced Engineering and Logistics, Samara, Russia
lebedev.kirill.oleg@gmail.com

Abstract

Freight transportation cost models have absorbed several decades of work on quantifying the external costs vehicles impose on other road users. A parallel literature has shown that compliance with duty-hour regulation is a measurably imperfect proxy for a driver's physiological fatigue state. These two bodies of work have not yet been combined. This article argues that the gap between them is now a pricing problem: multimodal physiological sensing has reached detection latencies and discrimination accuracy sufficient to generate a continuous, driver-specific impairment signal, yet freight cost models do not currently treat that signal as a priceable input, and externality-costing frameworks do not condition their per-mile crash-cost estimates on driver state. The article traces this omission through vehicle-routing, accident-externality, and driver-monitoring literatures. Measurement capability already exists; the absence of an institutional channel connecting a real-time physiological signal to a cost or liability mechanism is what blocks impairment pricing. The analysis concludes with the specific data linkage this channel would require and the regulatory gap currently preventing it.

Keywords: Driver fatigue, externality pricing, freight transportation economics, hours-of-service regulation, physiological monitoring, vehicle routing

Licensed under Creative Common

INTRODUCTION

A tabu-search comparison of scheduling algorithms for long-haul truck routing found that biomathematical fatigue constraints, derived from sleep-physiology models of alertness decay, generated measurably lower fatigue-risk exposure than schedules built purely to satisfy hours-of-service thresholds. Duty-hour accounting treats every hour within the legal window as equivalently safe regardless of the driver's accumulated sleep debt or circadian phase (Fu & Ma, 2022). This finding closes a question the routing literature had left open for over a decade: whether the eleven-hour driving window and ten-hour rest requirement that structure nearly every U.S. long-haul optimization model function as a genuine safety constraint or as an administrative convenience mistaken for one. They function as the latter. Every freight cost model built on top of HOS-compliant routing inherits the same substitution, treating regulatory compliance as if it were the underlying risk variable the cost function is meant to price.

The accident-externality literature has independently established that road transportation imposes costs on third parties that per-mile insurance pricing does not fully capture: traffic-dense states generate external accident costs of $1,725 to $3,239 per additional driver annually (Edlin & Karaca-Mandic, 2006), and a single shale-gas truck operating in a given zip code for a year raises neighboring drivers' insurance premiums by an estimated $0.48 through elevated collision risk alone (Muehlenbachs, Staubli, & Chu, 2021). Both estimates are computed as functions of vehicle miles traveled or truck presence, holding driver state constant, because no dataset available to either study measured driver state as a time-varying covariate. Multimodal physiological monitoring has since closed that measurement gap at the vehicle level: a fused five-channel Adaptive Physiological Impairment Monitoring System (APIMS), evaluated on 1,247 simulator sessions, achieved an area-under-curve of 0.973 and a mean detection latency of 1.77 seconds for graduated impairment classification (Doskach, 2026a), though validation to date rests on simulator data, a limitation the Literature Review returns to. The externality-pricing literature confines physiological signals like this one to vehicle-level safety engineering; the freight cost-modeling literature on which macroeconomic transmission analyses such as Doskach's (2026b) are built does the same. Doskach's two studies supply the components of such a connection, a validated sensing system and a macroeconomic cost-transmission model, and this article builds on them to trace why the connection has not yet been made and what would be required to make it.

LITERATURE REVIEW

Freight transportation cost research has become increasingly precise about which operational inefficiencies translate into macroeconomic cost, without extending that precision to driver physiological state. Hubbard (2003) demonstrated that pre-digital information frictions,

Licensed under Creative Common

specifically dispatchers' inability to observe truck location and available capacity in real time, constituted the primary structural cause of elevated deadhead mileage, and that onboard computing investment reduced deadhead by closing that observation gap. The mechanism is informational: the trucks and routes available for backhaul did not change, only the dispatcher's ability to see them. Doskach (2026b) extends this transmission logic to the macroeconomic level, modeling how carrier-level dispatch efficiency, deadhead reduction, and delivery reliability propagate into supply-chain cost and safety-stock requirements through a stochastic routing architecture that treats demand uncertainty via scenario sampling.

Hummels (2007) generalizes the same observational logic to international trade, showing that falling ocean and air freight rates raised trade volume less through the price mechanism itself than through the reduction in transit-time variance that better-tracked shipments made possible, since buyers priced schedule reliability as a cost distinct from the freight rate. The mechanism he isolates, that unobserved variance imposes a cost even when the mean outcome is acceptable, recurs inside freight-cost modeling at a smaller scale: a carrier's average on-time performance can satisfy a shipper's contractual threshold while the variance driving occasional severe delays stays invisible to any metric built from averaged delivery data. Doskach's (2026b) model is well placed to take up this logic, since its scenario-based architecture already treats delivery-reliability variance explicitly and can be extended to decompose that variance into its constituent sources. The observational gap Hummels traces at the scale of international trade statistics recurs, in miniature, inside a single carrier's dispatch data, and the signal described in Doskach (2026a) offers a way to close it.

Externality-pricing research supplies the methodological apparatus for converting an observed risk factor into a monetary cost, and its own internal evolution shows that apparatus can absorb new covariates once those covariates are measured. Edlin and Karaca-Mandic (2006) established that panel data on state-level insurance premiums and loss costs can recover a per-driver external cost estimate large enough to justify a corrective tax exceeding $220 billion annually nationwide, a result that depends on traffic density as the conditioning variable because density was the covariate available in the insurance data used. Muehlenbachs, Staubli, and Chu (2021) replicated this approach for trucking specifically. They used shale-gas development as a source of exogenous variation in truck traffic to isolate the accident externality trucks impose on surrounding car traffic, net of any general road-condition confound, a design that answers the identification problem Edlin and Karaca-Mandic's cross-sectional approach could not fully resolve. Dai, Yu, Liu, Cui, and Levinson (2025) push the same decomposition to finer geographic resolution, separating internal costs borne by crash-involved drivers from external costs imposed on other parties at the level of individual road links using safety performance

Licensed under Creative Common

functions and ordered probit severity models, and finding that crash costs remain too small relative to the value of travel time to influence route choice even where drivers are informed of them. Each successive study narrows the unit of analysis, from state to corridor to link, without introducing driver physiological state as a conditioning variable, because the datasets underlying them did not contain one. The methodological apparatus for pricing a newly measured risk factor is already built and already validated at increasingly fine spatial resolution. It is waiting on a covariate the crash-record data these studies use does not currently include.

Santos, Behrendt, Maconi, Shirvani, and Teytelboym (2010) situate the accident externality within a broader taxonomy of road-transport externalities that also includes congestion, local air pollution, noise, and infrastructure wear, and their review makes a point the narrower accident-costing literature tends to obscure: externalities differ not only in magnitude but in the observability of the variable that generates them. Congestion pricing succeeds because location and time of travel are directly observable through existing infrastructure such as toll gantries or cordon cameras. The accident externality has no equivalent observable proxy for the variable that actually generates risk, driver physiological state at the moment of travel, which is why Edlin and Karaca-Mandic, Muehlenbachs, Staubli, and Chu, and Dai, Yu, Liu, Cui, and Levinson each substitute an observable correlate, traffic density, truck presence, road-link characteristics, for the unobservable variable the cost is nominally meant to price. Santos et al.'s taxonomy clarifies that this substitution is a structural feature of accident-externality pricing as a category, and it explains why the accident-cost literature has converged on finer geographic resolution as its axis of methodological progress instead of finer driver-state resolution.

VEHICLE-ROUTING RESEARCH THAT SUBSTITUTES PREDICTION FOR MEASUREMENT

A second research stream has approached the same substitution problem from the scheduling side, and its progression documents exactly where prediction has advanced without being matched by measurement. Goel and Vidal (2014) built an optimization-based international assessment comparing HOS regimes in the United States, Canada, and the European Union, embedding each jurisdiction's duty-hour thresholds directly as hard constraints in a cost-minimizing routing algorithm. The approach takes compliance itself as the safety objective, which holds only if compliance and fatigue risk move together, an assumption the paper does not test. Zanella, Palma Lima, Mulati, Cardoza Galdamez, and Lapasini Leal (2025), synthesizing forty-one papers on vehicle routing under HOS regulation published between 2000 and 2024, confirm that this compliance-as-objective framing remains the dominant paradigm across the field, with the large majority of reviewed studies encoding legal duty-hour limits as

Licensed under Creative Common

the binding safety constraint. Fu and Ma (2022) is the documented exception within that same body of work: replacing HOS thresholds with biomathematical fatigue constraints derived from sleep-physiology models, and finding through tabu-search computation on both modified benchmark instances and a real-world case that schedules built around predicted alertness state consistently reduced fatigue-risk exposure relative to schedules built around duty-hour compliance alone. HOS constraints still bound duty hours to a legally enforceable limit; they remain a poor proxy specifically for the risk variable the routing objective is meant to minimize. That gap is why substituting a predicted physiological state for a compliance threshold produced a measurable improvement even without any sensor data collected from an actual driver.

Hanowski, Hickman, Fumero, Olson, and Dingus (2009) supply the naturalistic evidence the compliance-as-safety substitution in vehicle-routing models has long lacked. Instrumenting commercial vehicles under the 2003 hours-of-service revision with in-cab video and sleep monitoring, they found that drivers fully compliant with the revised duty-hour limits nonetheless obtained substantially less sleep than the regulation's biomathematical justification assumed, with a meaningful share of drivers averaging under five hours of sleep per rest period despite operating entirely within the legal duty window. The finding matters for the routing literature because it was collected from drivers already operating under the regime that Goel and Vidal (2014), and the forty-one studies Zanella, Palma Lima, Mulati, Cardoza Galdamez, and Lapasini Leal (2025) reviewed, encode as their binding safety constraint. Hanowski et al. lacked a continuous physiological impairment classifier of the kind APIMS later provided; their instrumentation measured sleep duration and video-coded drowsiness, not a fused real-time signal. Their data nonetheless demonstrates that the gap Fu and Ma (2022) later closed computationally, between HOS compliance and actual physiological state, is an empirically observed feature of the regulated driver population, measured well before a sensor able to track it continuously in revenue service existed.

Prediction is not the same instrument as measurement, and the gap between them is the one this article's core claim turns on. A biomathematical model like the one Fu and Ma (2022) embed generates an expected fatigue trajectory from a driver's scheduled duty pattern. Its output degrades under exactly the conditions freight operations most need it to hold: irregular start times, split sleep periods, and any deviation from the modeled sleep-wake assumption that a real driver's actual physiology may or may not follow. Williamson, Lombardi, Folkard, Stutts, Courtney, and Connor (2011) document the underlying chronobiological mechanism this dependency exposes. Fatigue intensity is jointly determined by homeostatic sleep pressure, circadian phase, and cumulative sleep history; none of these are captured by a duty-hour schedule, and none are captured at the individual level by a population-level biomathematical model. A driver who has

Licensed under Creative Common

accumulated sleep debt across several short rest periods can be fully compliant with both HOS rules and a biomathematically optimized schedule while operating at a physiological impairment level a scheduled prediction has no mechanism to detect. Closing this residual gap requires a signal generated by the driver's actual physiological state during the actual trip, not a signal generated by the schedule the driver was assigned before the trip began.

Driver-monitoring research has spent the past decade building toward exactly that signal, and its own internal limitations document why deployment has lagged behind laboratory capability. Vision-based systems monitoring eyelid closure, blink rate, and head pose represent the dominant commercially deployed category. They achieve area-under-curve values between 0.84 and 0.92 for binary drowsy-versus-alert classification in laboratory evaluation, but this discrimination degrades materially under variable illumination, with eyeglass-wearing operators, and across ethnically diverse populations for whom commercial facial-analysis models exhibit reduced sensitivity (Deng, Wu, & Yang, 2019). The most widely deployed monitoring category is therefore least reliable under exactly the lighting and demographic conditions a national trucking fleet actually operates in. Physiological alternatives address the accuracy problem at its source: electrodermal activity, photoplethysmography, and heart-rate variability have achieved positive predictive values exceeding 90 percent in laboratory protocols (Kaplan, Guvensan, Yavuz, & Karalurt, 2015; Vicente, Laguna, Bartra, & Bailón, 2016), and electroencephalographic measures detect fatigue onset preceding any observable behavioral change, with frontal theta and alpha power rising monotonically with sleep-deprivation duration (Ma et al., 2022). A systematic review of 387 publications on physiological driver-monitoring found that hybrid systems fusing physiological and behavioral signal classes outperformed single-modality baselines by 0.05 to 0.10 AUC units on average. The authors attribute the gain to the complementary time scales of different signal types: autonomic measures like EDA and heart-rate variability unfold over minutes, eyelid kinematics and head motion capture state transitions over seconds, and steering telemetry integrates behavioral consequences over tens of seconds (Liang, Liu, Dong, & Pang, 2022). Their sample was dominated by two-signal fusion architectures; whether adding a third, fourth, or fifth modality continues to yield positive returns is not a question their data answers, though the timescale-complementarity mechanism they propose would predict diminishing but still positive returns from doing so, an extrapolation offered here as the article's own inference.

Ramzan, Khan, Awan, Ismail, Ilyas, and Mahmood's (2019) survey of drowsiness-detection techniques maps the field across three input classes, behavioral, physiological, and vehicular, and documents a pattern consistent with the fusion logic Liang, Liu, Dong, and Pang (2022) advance: single-modality systems within each class had by 2019 already approached a performance ceiling that further refinement within the same modality was unlikely to break. The

Licensed under Creative Common

survey's own conclusion, that further gains would require combining input classes, anticipates the fusion architectures that followed. Zhang, Lyu, and Qin (2023) test that prediction directly with a mixed-attention multimodal fusion model that learns to weight eye, mouth, and head-pose features dynamically, reporting accuracy gains over static-weighting baselines the authors attribute to the model's ability to down-weight a degraded channel, such as an eye-region feature obscured by sunglasses,. The mechanism differs architecturally from APIMS's temporal-convolutional and bidirectional-LSTM fusion of physiological and behavioral channels, but the lesson is the same: treating input channels as equally reliable at every moment discards information a learned weighting scheme can recover. Neither study addresses the externality-pricing question this article raises; both operate entirely within the driver-monitoring literature's own objective of maximizing detection accuracy.

None of these physiological modalities individually solved the deployment problem. Sikander and Anwar's (2019) scoping review of driver-monitoring systems found that commercial trucking deployments remained uniformly limited to single-modality camera-based drowsiness alerts, with no deployed system integrating more than one physiological channel or coupling detection to any automated vehicle response. Laboratory-validated technology existed well before any commercial system combined it into a passively acquired, multimodal instrument. Doskach's (2026a) evaluation of APIMS closes precisely that combination gap. The system fuses five passively acquired channels, infrared facial video, steering-wheel electrodermal sensing, headrest-embedded photoplethysmography, inertial head kinematics, and CAN-bus steering telemetry, into a hybrid temporal-convolutional and bidirectional-LSTM classifier. Streams are synchronized to a shared 2 Hz feature epoch; each epoch is assigned to one of five graduated impairment states instead of a binary alert-or-not output. On a held-out validation fold of 1,247 simulator sessions, this five-way fusion achieved an area-under-curve of 0.973, a sensitivity of 94.7 percent at 96.8 percent specificity, and a mean detection latency of 1.77 seconds, an improvement of 0.082 AUC units over the best individual modality (infrared camera, AUC 0.891). That margin sits at the upper edge of the fusion advantage Liang et al. document across the wider literature, extending it beyond the two-signal architectures their sample was built from. The system's discrimination of a distinct pharmacological-impairment class extends the externality-pricing apparatus reviewed above most directly: the class is identifiable through the simultaneous combination of tonic electrodermal elevation, reduced heart-rate variability, and oculomotor slowing without the head-nodding pattern camera-only systems are trained to detect, at 88.4 percent sensitivity and 97.1 percent specificity. Pharmacological impairment from legally obtained sedating medication is a risk category HOS accounting cannot address by construction, since it is orthogonal to duty-hour accumulation, and it is a risk category the crash-record data underlying

Licensed under Creative Common

Edlin and Karaca-Mandic (2006), Muehlenbachs, Staubli, and Chu (2021), and Dai, Yu, Liu, Cui, and Levinson (2025) has not separately identified. A driver-state signal now exists, at least under simulator conditions, that could condition an externality estimate the way traffic density or shale-gas well proximity has conditioned prior estimates. Externality studies have not yet used it. The sensing infrastructure and the crash-costing infrastructure were built by separate research communities working from separate datasets.

A measured signal and a costing methodology are not sufficient if the regulatory and enforcement apparatus meant to translate risk into consequence cannot reliably observe the carriers generating the risk. Table 1 demonstrates that the U.S. commercial-motor-vehicle safety system already generates large volumes of inspection and compliance data, yet these data remain episodic and administrative in nature. They do not provide a continuous, driver-specific measure of physiological impairment that could be incorporated into crash-risk estimation or external-cost pricing.

Table 1: Scale of U.S. Large-Truck Operations and Limitations of Current Safety Monitoring

Indicator	**Value**	**Reference period**
Vehicle miles traveled by large trucks	327.0 billion vehicle-miles	2021
Active regulated motor carriers	813,844	2022
Carriers operating one power unit	445,673	2022
Total commercial motor vehicle inspections	2,984,331	2022
Driver inspections	2,879,966	2022
Driver out-of-service rate	6.9%	2022
Vehicle inspections	2,016,103	2022
Vehicle out-of-service rate	22.6%	2022
“False report of driver’s record of duty status” violations	60,856	2022
Electronic logging device (ELD) record-of-duty-status violations	40,061	2022
“Record of duty status not current” violations	22,374	2022
Estimated crashes in which driver drowsiness was a contributing factor	8.8–9.5% of 701 studied crashes	Naturalistic-driving study
Estimated drowsiness contribution in severe crashes	10.6–10.8%	Naturalistic-driving study
Estimated share of fatal crashes involving drowsiness	Up to 21%	2009–2013 estimate

Source: FMCSA (2023) for vehicle miles traveled, carrier counts, inspections, out-of-service rates, and violations; Owens et al. (2018) for the naturalistic-driving drowsiness estimates; Tefft (2014) for the share of fatal crashes.

Licensed under Creative Common

The Federal Motor Carrier Safety Administration's Compliance, Safety, Accountability program aggregates roadside inspection and crash data into carrier-level percentile scores intended to prioritize enforcement. A preliminary investigation of FMCSA inspection violations found that this aggregation obscures whether the violations driving a given carrier's crash-risk correlation cluster within specific vehicles, drivers, or operational periods (Lack, Berkow, & Gao, 2021). A carrier-level score cannot distinguish a carrier with one severely impaired driver from a carrier with diffuse, low-severity noncompliance across its fleet. Knipling, Hyten, and Singer (2011) compound this resolution problem at the population level, finding that carriers operating the majority of U.S. commercial truck-miles are systematically subject to lower inspection frequency and therefore generate less statistically reliable CSA scores, despite representing a disproportionate share of at-fault crashes. Kim, Park, and Wagner (2024) sharpen this institutional gap into a specific finding about the population most exposed to it: owner-operator experience moderates violation outcomes. Less experienced owner-operators, who by Hubbard's (2003) transaction-cost logic are the entrants least able to absorb the fixed cost of continuous fleet-level monitoring infrastructure, carry both the least reliable CSA data and, plausibly, the least institutional capacity to detect their own physiological impairment through any channel other than a passively worn sensor. Cantor, Corsi, Grimm, and Özpolat (2010) build driver-level covariates such as violation history and tenure directly into a crash prediction model, precisely because carrier-level aggregation of the kind Lack, Berkow, and Gao (2021) critique loses the individual-driver signal. A physiological monitor would restore that signal directly, without requiring the roadside-inspection sampling process that currently generates CSA data to observe the driver at all.

The National Academies of Sciences, Engineering, and Medicine's (2012) synthesis of safety management practice among small motor carriers reaches a conclusion consistent with the informational argument developed above. Surveying carrier safety programs directly, the synthesis found that small carriers frequently possessed safety knowledge comparable to larger fleets but lacked any mechanism for converting that knowledge into a continuous record of driver state, relying instead on periodic supervisor observation and self-report. The synthesis frames this as a resource gap, consistent with the small-carrier-capacity account this article contests. Read alongside Lack, Berkow, and Gao's (2021) finding that CSA aggregation obscures whether risk concentrates in specific drivers, the National Academies data supports a narrower reading: the missing element was not safety knowledge or organizational commitment but a low-cost, passively acquired instrument capable of generating the continuous record manual practices could not produce. A large carrier with a dedicated safety department faces the identical instrument gap; supervisor observation and self-report do not scale to continuous

Licensed under Creative Common

coverage regardless of headcount, which only affords more staff to attempt the manual approximation the synthesis documents as structurally inadequate at any scale.

DISCUSSION

The literature review exposes a tension with a specific shape. Impairment can now be measured, at sub-two-second latency and above-0.97 discrimination, at least in simulator conditions (Doskach, 2026a). Externalities can be priced, through three independent methods spanning nearly two decades that have progressively narrowed their unit of analysis from state to corridor to individual road link (Edlin & Karaca-Mandic, 2006; Muehlenbachs, Staubli, & Chu, 2021; Dai, Yu, Liu, Cui, & Levinson, 2025). Measurement and pricing capability both exist; what connects a real-time physiological signal to a pricing or enforcement mechanism does not. The CSA scoring apparatus reviewed above is built to aggregate roadside inspection snapshots, not to ingest continuous vehicle telemetry. The externality-costing literature has built its crash-cost estimates from insurance and crash-record panels that were not designed to accept a driver-state covariate. Figure 1 presents the proposed architecture for linking a continuous, driver-specific physiological impairment signal to crash-risk estimation, external-cost valuation, and institutional risk internalisation.

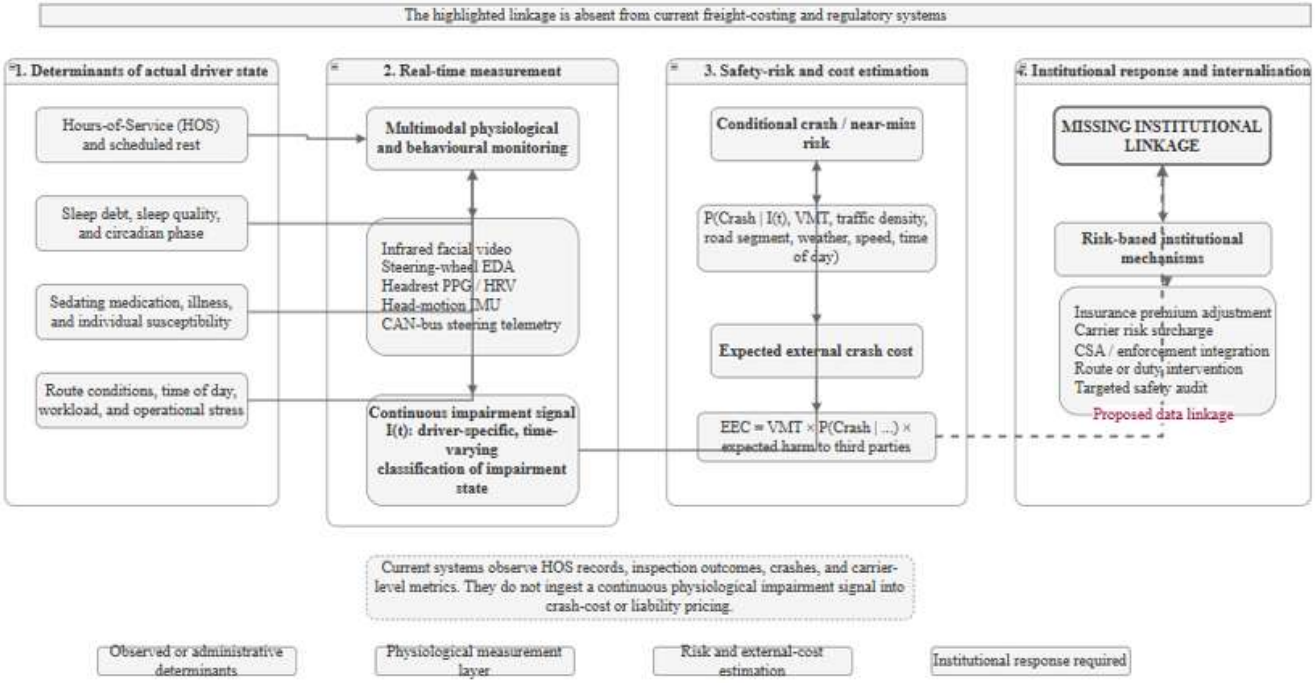


Figure 1: Proposed architecture for incorporating driver physiological impairment into freight crash-externality pricing. The real-time measurement layer is based on the APIMS of Doskach (2026a); the proposed data linkage connects its impairment signal to crash-cost estimation.

Licensed under Creative Common

The gap identified above also reframes what a driver-state signal would need to interoperate with institutionally. CSA's percentile architecture is built around discrete inspection events; the Doskach (2026a) classifier and comparable multimodal systems generate a continuous stream of graduated impairment states at a two-hertz epoch. Ingesting a continuous signal into an architecture built for episodic events requires deciding what threshold along a continuous impairment gradient triggers a discrete regulatory or pricing consequence, a decision current CSA methodology has not had to make because its existing inputs do not vary continuously. The externality-pricing studies reviewed above face a parallel translation problem from the other direction: their crash-cost estimates are computed from discrete crash and near-miss events, and converting a continuous impairment stream into a covariate for those models requires an aggregation rule, whether the relevant unit is peak impairment during a trip, cumulative time above a threshold, or a duration-weighted exposure measure, that the studies cited above have not specified because they have not had continuous driver-state data to aggregate. Both translation problems are solvable in principle; these questions have not yet been posed in either literature, because posing them requires a research program situated across both fields simultaneously.

The small-carrier safety literature attributes the persistence of this gap to organizational capacity, specifically a lack of safety-department resources or safety-culture maturity among small carriers (Knipling, Hyten, & Singer, 2011). The evidence assembled in this review points to an informational constraint instead. A well-resourced carrier with a mature safety department gains no additional ability to price driver impairment under the current CSA architecture, because the architecture itself has no channel for ingesting a continuous physiological signal regardless of the carrier's internal capacity to generate one. Lack, Berkow, and Gao's (2021) finding that carrier-level violation aggregation cannot distinguish concentrated from diffuse risk within a fleet holds for a large, well-resourced carrier exactly as it holds for a small one, because the limitation sits in the scoring methodology's data structure, not in any carrier's compliance effort. If organizational capacity were the binding constraint, larger carriers with dedicated safety staff should already show measurably better-calibrated risk pricing than small carriers under the existing CSA framework. The aggregation problem Lack, Berkow, and Gao document is a property of the scoring architecture itself, and it constrains carriers of every size equally, a pattern consistent with an informational bottleneck and inconsistent with a story in which small-carrier under-resourcing is the primary obstacle to impairment pricing.

A second, more localized point concerns how the signal could be integrated into the macroeconomic framework of Doskach (2026b). The model attributes delivery-reliability degradation, the time-window violations that propagate into supply-chain safety-stock cost, to

Licensed under Creative Common

residual variance in demand forecasting once observable operational factors are controlled for. This attribution is well suited to the portion of reliability degradation the model's architecture can observe, since its scenario-sampling approach to demand uncertainty fits genuine forecast error. Williamson et al.'s (2011) chronobiological evidence establishes that a driver can be fully compliant with every observable scheduling constraint such a model tracks, hours, routes, delivery windows, while operating under an accumulated sleep debt. Fu and Ma (2022) demonstrated that HOS-compliant schedules and physiologically optimized schedules produce measurably different fatigue-risk profiles even when both satisfy identical duty-hour constraints. Some portion of the reliability degradation assigned to demand-forecast error may therefore also reflect driver-state degradation compounding across an extended duty cycle. The scope of the model followed the data available to it, which did not include driver state. The APIMS classifier of Doskach (2026a) provides a candidate variable for exactly this purpose: adding it as a covariate would allow the residual term to be decomposed into demand and driver-physiological components, which is the linkage the preceding paragraph identifies as the next step for the field.

CONCLUSION

The sensing technology capable of isolating the physiological component, the APIMS developed by Doskach (2026a), operates at sub-two-second detection latency with area-under-curve discrimination above 0.97, a level of technical maturity established under simulator conditions and requiring further validation under road-recorded fleet operation before its output can serve as a production-grade input to a cost model. The task that remains is a data-linkage task: construction of a pipeline that matches continuous physiological classification output against crash and near-miss records at the same geographic and temporal resolution already used in link-level externality-costing frameworks. Completion and validation of this linkage against operational fleet data constitutes the specific empirical step required to convert the externality identified in this analysis into a priced, actionable input within freight transportation cost models.

The architecture required for this linkage carries a further implication for the function an externality estimate performs once conditioned on driver physiological state. Established externality-pricing methods quantify a cost generated by presence in traffic, a cost that exists independently of any individual driver's behavior at a given moment. A physiological impairment signal quantifies a cost tied to a specific driver's state during a specific interval, so incorporating this signal transforms a population-level externality estimate into an individual, moment-specific liability signal, carrying direct consequences for insurance underwriting, carrier dispatch liability,

Licensed under Creative Common

and driver privacy. Owner-operators and other capital-constrained carriers currently generate the least reliable safety data under prevailing inspection-based scoring systems and would constitute the population first exposed to physiological pricing, precisely because aggregate risk metrics protect this population least. The institutional channel through which this linkage is implemented sets its social function. A regulatory mandate anchored in safety compliance directs the technology toward crash reduction across the driver population as a whole. An underwriting or liability channel directs the same technology toward differentiated risk pricing that concentrates cost on the least capitalized carriers. Institutional design therefore determines the ultimate social value of the technical capability documented in this analysis.

Several U.S. states already permit usage-based auto insurance pricing conditioned on telematics data unrelated to physiological state, establishing legal and actuarial precedent for individual-level risk pricing in a closely related domain before any physiological covariate has entered the freight-crash-costing pipeline reviewed here. Extending that precedent to a physiological impairment covariate would not require new legislative authority so much as a decision, currently absent from FMCSA rulemaking and from the actuarial literature alike, about which institutional actor is authorized to receive a continuous impairment stream and for what purpose. Absent that decision, the technical capability documented across the driver-monitoring literature, including Doskach's APIMS, will continue to accumulate laboratory validation without a regulatory or contractual channel through which it can reach the crash-costing frameworks this article has shown are otherwise ready to receive it.

Several lines of further research follow from this analysis. The first is on-road validation of the sensing layer. Because the APIMS results rest on simulator sessions, detection latency and discrimination should be re-estimated on road-recorded fleet data and stratified by illumination, eyewear, and driver demographics, the conditions under which the vision-based systems reviewed above degrade. The second concerns the aggregation rule that converts a continuous impairment stream into a covariate for crash-cost models. Peak impairment during a trip, cumulative time above a threshold, and duration-weighted exposure can be compared directly against crash and near-miss outcomes at link-level resolution, and the comparison would show which specification most improves externality estimates. The third is a test that separates the informational-constraint account from the organizational-capacity account. If capacity were the binding constraint, larger carriers with dedicated safety staff should already show better-calibrated risk pricing under CSA. The fourth is institutional: pilot comparisons of a safety-compliance channel and an underwriting or liability channel could measure how each distributes cost across carriers of different capitalization, with owner-operators as the population of primary interest, and which privacy safeguards each requires. A related technical question,

Licensed under Creative Common

whether a third, fourth, or fifth signal class continues to improve discrimination beyond the two-signal architectures that dominate the fusion literature, can be tested within the same validation program.

REFERENCES

Cantor, D. E., Corsi, T. M., Grimm, C. M., & Özpolat, K. (2010). A driver focused truck crash prediction model. *Transportation Research Part E: Logistics and Transportation Review*, *46*(5), 683–692.

Dai, S., Yu, L., Liu, Z., Cui, M., & Levinson, D. (2025). The internal and external cost of motor vehicle crashes. *Scientific Reports*, *15*, 5441. https://doi.org/10.1038/s41598-025-89058-1

Doskach, H. (2026a). *Safety and compliance in commercial transportation.* LAP LAMBERT Academic Publishing.

Doskach, H. (2026b). Supply chain interdependencies and the macroeconomic role of surface freight transportation: Evidence from a joint probabilistic forecasting and stochastic routing architecture. *European Science Review*, (3–4), 74–78. https://doi.org/10.29013/ESR-26-3.4-74-78

Edlin, A. S., & Karaca-Mandic, P. (2006). The accident externality from driving. *Journal of Political Economy*, *114*(5), 931–955. https://doi.org/10.1086/508030

Federal Motor Carrier Safety Administration. (2023). *2023 pocket guide to large truck and bus statistics* (Report No. RRA-23-003). U.S. Department of Transportation.

Fu, J., & Ma, L. (2022). Long-haul vehicle routing and scheduling with biomathematical fatigue constraints. *Transportation Science*, *56*(2), 404–435. https://doi.org/10.1287/trsc.2021.1089

Goel, A., & Vidal, T. (2014). Hours of service regulations in road freight transport: An optimization-based international assessment. *Transportation Science*, *48*(3), 391–412.

Hanowski, R. J., Hickman, J., Fumero, M. C., Olson, R. L., & Dingus, T. A. (2009). The sleep of commercial vehicle drivers under the 2003 hours-of-service regulations. *Accident Analysis and Prevention*, *41*(5), 975–985. https://doi.org/10.1016/j.aap.2009.05.008

Hubbard, T. N. (2003). Information, decisions, and productivity: On-board computers and capacity utilization in trucking. *American Economic Review*, *93*(4), 1328–1353. https://doi.org/10.1257/000282803769206322

Hummels, D. (2007). Transportation costs and international trade in the second era of globalization. *Journal of Economic Perspectives*, *21*(3), 131–154. https://doi.org/10.1257/jep.21.3.131

Kaplan, S., Guvensan, M. A., Yavuz, A. G., & Karalurt, Y. (2015). Driver behavior analysis for safe driving: A survey. *IEEE Transactions on Intelligent Transportation Systems*, *16*(6), 3017–3032. https://doi.org/10.1109/TITS.2015.2462084

Kim, S., Park, A., & Wagner, S. M. (2024). Owner-operator experience and violations in trucking: The mixed role of perceived justice. *Transportation Research Part E: Logistics and Transportation Review*, *188*, 103642. https://doi.org/10.1016/j.tre.2024.103642

Lack, C. D., Berkow, K. S., & Gao, Y. (2021). Insights into motor carrier crashes: A preliminary investigation of FMCSA inspection violations. *Accident Analysis & Prevention*, *155*, 106105. https://doi.org/10.1016/j.aap.2021.106105

Liang, Z., Liu, G., Dong, J., & Pang, M. (2022). A systematic review of physiological signals based driver drowsiness detection methods. *Sensors*, *22*(21), 8441. https://doi.org/10.3390/s22218441

Ma, Y., Chen, B., Li, R., Wang, C., Wang, J., She, Q., Luo, Z., & Zhang, Y. (2022). The application of electroencephalogram in driving safety: Current status and future prospects. *Frontiers in Psychology*, *13*, 919695. https://doi.org/10.3389/fpsyg.2022.919695

Muehlenbachs, L., Staubli, S., & Chu, Z. (2021). The accident externality from trucking: Evidence from shale gas development. *Regional Science and Urban Economics*, *88*, 103630. https://doi.org/10.1016/j.regsciurbeco.2020.103630

National Academies of Sciences, Engineering, and Medicine. (2012). *Safety management in small motor carriers* (CTBSSP Synthesis 22). Transportation Research Board. https://doi.org/10.17226/14637

Licensed under Creative Common

Owens, J. M., Dingus, T. A., Guo, F., Fang, Y., Perez, M., McClafferty, J., & Tefft, B. C. (2018). *Prevalence of drowsy-driving crashes: Estimates from a large-scale naturalistic driving study* (Research Brief). AAA Foundation for Traffic Safety.

Ramzan, M., Khan, H. U., Awan, S. M., Ismail, A., Ilyas, M., & Mahmood, A. (2019). A survey on state-of-the-art drowsiness detection techniques. *IEEE Access*, *7*, 61904–61919. https://doi.org/10.1109/ACCESS.2019.2914373

Santos, G., Behrendt, H., Maconi, L., Shirvani, T., & Teytelboym, A. (2010). Part I: Externalities and economic policies in road transport. *Research in Transportation Economics*, *28*(1), 2–45.

Sikander, G., & Anwar, S. (2019). Driver fatigue detection systems: A review. *IEEE Transactions on Intelligent Transportation Systems*, *20*(6), 2339–2352. https://doi.org/10.1109/TITS.2018.2868499

Tefft, B. C. (2014). *Prevalence of motor vehicle crashes involving drowsy drivers, United States, 2009-2013* (Technical Report). AAA Foundation for Traffic Safety.

Vicente, J., Laguna, P., Bartra, A., & Bailón, R. (2016). Drowsiness detection using heart rate variability. *Medical and Biological Engineering and Computing*, *54*(6), 927–937. https://doi.org/10.1007/s11517-015-1448-7

Williamson, A., Lombardi, D. A., Folkard, S., Stutts, J., Courtney, T. K., & Connor, J. L. (2011). The link between fatigue and safety. *Accident Analysis and Prevention*, *43*(2), 498–515. https://doi.org/10.1016/j.aap.2009.11.011

Zanella, A. F., Palma Lima, R. H., Mulati, M. H., Cardoza Galdamez, E. V., & Lapasini Leal, G. C. (2025). Vehicle routing and scheduling under hours of service regulations: A review. *Transportation Research Part A: Policy and Practice*, *201*, 104665. https://doi.org/10.1016/j.tra.2025.104665

Zhang, X., Lyu, Y., & Qin, L. (2023). Driver drowsiness detection using mixed-attention mechanism multi-modal fusion model. *Sensors*, *23*(3), 1408. https://doi.org/10.3390/s23031408

Licensed under Creative Common