Dynamic CMV Driver Scheduling Optimization Model under HOS Constraints

Li Zhao, Ph.D.[*]
Department of Civil and Environmental Engineering, University of Nebraska-Lincoln
2200 Vine Street, Lincoln, Nebraska, 68516
ORCID: 0000-0003-1914-1041
Email: lizhao@unl.edu

Hania Afzal, Ph.D.
Department of Civil and Environmental Engineering
2200 Vine Street, Lincoln, Nebraska, 68516
Email: hafzal2@unl.edu

Jiahe Cao
School of Computing, University of Nebraska-Lincoln
262 Avery Hall, 1144 T St, Lincoln, Nebraska, 68588
Email: jcao10@huskers.unl.edu

Nathan Huynh, Ph.D.
Department of Civil and Environmental Engineering, University of Nebraska-Lincoln
2200 Vine Street, Lincoln, Nebraska, 68516
Email: nathan.huynh@unl.edu

Jason Hawkins, Ph.D.
Department of Civil Engineering, University of Calgary
2500 University Drive NW, Calgary, T2N 1N4 Canada
Email: jfhawkin@ucalgary.ca

Total Number of Pages: 20

Submitted: July 31, 2026
[*]Corresponding author

ABSTRACT

Objectives: Commercial motor vehicle drivers must schedule their driving and rest under federal hours-of-service (HOS) limits, and on the road the right decision of stop or go is often not obvious given traffic conditions and parking availability at nearby rest areas. To address this, we propose a decision-support approach that plans a legal schedule before the trip and updates it during the trip. The objective of this study is to develop the scheduling framework underlying this approach, with the goal of reducing HOS violations without sacrificing delivery timeliness.

Methods: We use a mixed-integer program (MIP) to produce an HOS-compliant pre-trip plan on a time-space network, and a finite-horizon stochastic dynamic program (DP) that revises the baseline plan as travel delay and parking information become available. We treat HOS limits as soft constraints, so violations are penalized rather than forbidden. We solve the MIP using branch-and-cut and the DP using exact backward induction. We further use a 30-day electronic logging device record to reconstruct the four HOS clock states and a driving-simulator study of 12 CMV drivers across two controlled scenarios to obtain trip instances for validating the scheduling models.

Findings: Across trip length, parking availability, and delay frequency, the adaptive schedule holds the HOS violation rate low, compared with a baseline schedule that exceeds 90% as trips approach the driving limit. The benefit is largest for long trips, where the baseline almost always fails. Our results show that at 10 driving hours the adaptive schedule decreases violations (by inserting a required rest period) from 93% to 3.5%, but the same rest delays arrival, so on-time completion falls from 100% to 41.5%.

Novelty: The framework builds the HOS clock rules into both a pre-trip model that plans the baseline schedule and an in-trip adaptive model that updates it while the trip is underway. We demonstrate state reconstruction on real-world driving and simulator data for verification. To our knowledge, no prior work brings these parts together. This sets up the driver-facing guidance that current tools do not provide.

Practical Applications: Fleet managers can identify trips that cannot be driven legally before assigning them, and drivers can receive guidance on when to stop so they stay HOS-compliant without sacrificing delivery timeliness. Reconstructed clock states also let enforcement distinguish violations driven by unrealistic scheduling from those driven by individual driver behavior, which could inform enforcement or fleet safety review.

INTRODUCTION

Commercial motor vehicle (CMV) driver scheduling is a time-dependent decision problem that directly affects freight reliability, driver fatigue, truck parking demand, and highway safety. During a long-haul trip, a driver must decide when to drive, when to stop, where to stop, when to take a required break, when to begin a longer rest period, and whether to reroute when planned conditions change. These decisions are constrained by federal hours-of-service (HOS) rules by the Federal Motor Carrier Safety Administration (FMCSA), including the 11-hour driving limit, 14-hour duty window, 30-minute break requirement (after 8 hours of continuous driving), and 60/70-hour 7/8-day limit for property-carrying drivers (FMCSA 2022). Although these rules define the legal structure of CMV operations, they do not directly provide optimal driving, stopping, resting, and rerouting schedules under changing roadway conditions and delivery deadlines.

In practice, a driver can fall behind schedule that is feasible at the start of a trip once traffic congestion, work zones, adverse weather, delivery delays, or limited truck parking add time to the trip. These disruptions require drivers to make sequential, situation-dependent decisions based on their current location, remaining driving and duty-window time, time since the last break, parking availability, and delivery deadline. For example, continuing to drive may appear beneficial at one stage but create greater future risk if delays reduce the remaining legal driving time or parking becomes unavailable near the end of the duty window. CMV driver scheduling is therefore more complex than a static shortest-path or delivery-time minimization problem because each decision must account for both the driver's evolving HOS state and uncertain future network conditions.

Prior research has examined CMV fatigue and safety, HOS-constrained truck scheduling, truck parking availability, and dynamic vehicle routing (Soccolich et al. 2013; Goel and Irnich 2017; Pilz et al. 2022; Pillac et al. 2013). However, these issues have largely been addressed in separate research streams, despite the strong influence of working-hour regulations on route feasibility and schedule design (Goel and Kok 2012). This separation highlights the need for an integrated dynamic scheduling framework that jointly considers regulatory constraints, parking availability, and changing operating conditions.

This paper develops a two-stage optimization framework for CMV driver scheduling under HOS constraints. In the first stage, a mixed-integer programming model generates an initial schedule for the driver. In the second stage, when the driver is en route, a dynamic programming model updates the remaining schedule as operating conditions change. The framework explicitly considers the driver's state using location, accumulated driving time, prior duty history, parking availability, traffic conditions, weather, work zones, and delivery deadlines, and decisions are provided in the form of driving, performing on-duty non-driving activities, taking short non-driving breaks, and taking an extended rest period. The framework's basic behavior is examined using synthesized trip scenarios based on real-world CMV driving records and controlled driving-simulator data, which are used to evaluate stop-versus-proceed decisions under dynamic conditions.

RELATED WORK

Research relevant to CMV driver scheduling under HOS constraints has developed largely in isolation across several disparate areas, including driver fatigue and safety, HOS-constrained optimization, truck parking, and operational uncertainty. The following review is organized by category to show how it builds on and differs from prior work.

CMV Driver Fatigue, Work-Rest Patterns, and Safety

CMV driver scheduling is closely related to fatigue and safety because long-haul driving involves more than time spent behind the wheel. Drivers also spend time loading and unloading, waiting, inspecting vehicles, searching for parking, and handling other on-duty non-driving tasks. These activities can affect fatigue even when the driver remains within legal driving limits. Naturalistic driving research has shown that driving performance and safety-related outcomes are associated with driving duration, non-driving work, rest breaks, and work-rest patterns (Blanco et al. 2011; Soccolich et al. 2013). This means that schedule quality should not be measured only by total travel time or legal compliance. A schedule can be legal but still create high fatigue risk or additional financial costs if it includes long on-duty periods, irregular rest timing, limited sleep opportunity, or nighttime driving.

The safety literature also shows that the timing and structure of rest breaks matter. Studies of CMV crash risk have found that driving patterns, rest-break periods, and time-dependent work exposure are associated with safety outcomes (Jovanis et al. 2012; Chen and Xie 2014a,b). These findings support the use of schedule-level variables such as cumulative driving time, cumulative on-duty time, time since last break, and multi-day duty history. They also suggest that fatigue-related and safety-related costs should be represented in CMV scheduling models rather than treated as separate outcomes after a schedule has already been generated.

HOS-Constrained Truck Driver Scheduling

Truck driver scheduling under HOS rules has been widely studied in operations research. This body of work has shown that driving time, duty time, break requirements, and rest periods can be represented as mathematical constraints in truck scheduling and vehicle routing models. These constraints can substantially change feasible routes, delivery timing, and stop locations (Goel 2010; Goel and Kok 2012; Rancourt et al. 2013; Goel and Vidal 2014). A route that is optimal based only on distance or travel time may become infeasible once HOS rules are considered. Therefore, driver scheduling needs to be solved together with routing, delivery time windows, and stopping decisions (Goel and Irnich 2017).

Existing HOS-constrained models provide the foundation for the proposed pre-trip optimization model. They show how to track driving time since the last qualifying rest period, on-duty time within a duty window, break timing, and multi-day duty limits. They also show that exact algorithms, dynamic programming heuristics, and routing-and-scheduling formulations can be used to solve structured truck scheduling problems (Kok et al. 2010; Goel and Irnich 2017). These methods are useful for generating an initial trip plan when travel times, service times, and rest opportunities are known in advance.

However, CMV operations are often not fully deterministic. Time-dependent congestion and other roadway disruptions can change travel times and route costs, while truck-parking availability varies across locations and times (Kok et al. 2012; Pillac et al. 2013; Vital and Ioannou 2021). Because HOS clocks continue to run during delay, waiting, and parking search, even a modest change in travel time can affect the rest plan and delivery feasibility. A practical scheduling model should adjust the pre-trip plan and update the remaining schedule as operating conditions change.

Truck Parking and Stop-Location Decisions

Truck parking is a key element of CMV driver scheduling because breaks and rest periods require safe and legal places to stop. National assessments have reported persistent shortages of truck

parking spaces in the United States, especially along major freight corridors and near large urban areas (Federal Highway Administration 2020). Parking shortages can force drivers to stop earlier than planned, search for parking, reroute to another facility, wait for a space, or park in unauthorized locations (e.g., shoulder lane, off-ramps). Recent optimization studies have begun to integrate parking availability into truck scheduling models and shown that parking availability can change both the best route and the best rest plan (Vital and Ioannou 2019, 2021). Other work has examined coordinated assignment of parking spaces for breaks and rest periods and shown that parking management can be modeled as a system-level scheduling problem (Pilz et al. 2022). These studies treated parking as a limited resource that interacts with driver time regulations.

For CMV driver scheduling, parking availability should be represented as a time-dependent state rather than as a fixed attribute of a stop. Empirical truck-parking data show that the difficulty of finding parking varies by time of day and that drivers consider expected availability when selecting locations for required rest periods (Boris and Brewster 2018). A location that is feasible earlier in the day may become unavailable later, making the driver's best action dependent on expected conditions at downstream facilities. Because HOS requirements determine when the driver must stop, parking and stopping decisions should be evaluated jointly with the remaining driving time, duty-window time, and break status (Gaddy et al. 2018; Vital et al. 2021). Real-time parking detection and information systems further demonstrate that parking availability can be observed and updated dynamically during trip execution (Morris et al. 2024).

Cost Implications of HOS Constraints and Operational Uncertainty

HOS regulations affect not only the legal feasibility of truck schedules but also their operational and economic performance. Driving-time and duty-window limits may require earlier stopping, additional rest periods, route changes, or longer trip durations, thereby increasing travel time, driver cost, and the risk of delayed delivery. Previous optimization studies have shown that working-hour regulations can substantially alter route feasibility, schedule structure, and total transportation cost (Goel and Kok 2012; Rancourt et al. 2013; Goel and Vidal 2014; Goel and Irnich 2017). Consequently, a schedule that does not explicitly account for the operational costs induced by HOS constraints may underestimate the total trip cost required to complete a trip legally.

Operational uncertainty can further increase the cost of an initially feasible schedule. Congestion, adverse weather, work-zone disruptions, and service delays can increase travel and waiting times, reduce the remaining HOS margins, and create early- or late-delivery penalties (Kok et al. 2012; Taş et al. 2013, 2014). Because these conditions may change after departure, the schedule that minimizes cost at the beginning of a trip may no longer remain the lowest cost. Dynamic-routing research therefore supports revising the remaining route or schedule as new information becomes available (Pillac et al. 2013; Ritzinger et al. 2016; Psaraftis et al. 2016).

Research Gaps

The existing literature provides important building blocks but does not fully address dynamic CMV driver scheduling under HOS constraints. Safety studies show that fatigue and work-rest patterns matter, and an Electronic Logging Device (ELD) is now mandated to regulate the work-rest balance and can be used for measuring schedule patterns (Hickman et al. 2017; Crizzle et al. 2022). However, they are mainly retrospective. Truck parking studies show that parking availability affects stop-location decisions, but parking is not always linked with the full HOS clock structure. Dynamic routing models provide tools for online decision-making, but they often do not include

the specific HOS state variables that govern CMV driver schedules. This paper addresses these gaps through a two-stage optimization framework that integrates HOS constraints with CMV driver scheduling. The first stage generates an initial, HOS-compliant schedule using mixed-integer programming, providing the structured, forward-looking planning that ELD-based retrospective analysis cannot offer. The second stage employs a dynamic programming model that updates the remaining schedule en route as traffic conditions, weather, work zones, delivery delays, and parking availability change, directly linking parking and other time-dependent disruptions to the driver's evolving HOS state rather than treating them as separate considerations. By representing this set of state variables within a single adaptive framework, the proposed model extends dynamic routing approaches to the specific regulatory and operational structure of CMV driver scheduling.

METHODS

This section presents the mathematical formulation of the pre-trip and in-trip models. The pre-trip model determines a feasible driving-rest schedule before departure, while the in-trip model updates this schedule as the driver's HOS clock state evolves and real-time conditions are observed. Sets, parameters, and decision variables for both formulations are summarized in Table 1.

Time–Space Network and State

The problem is formulated on a time–space network at resolution $\Delta = 0.5$ h. A state is observed at each stage $t \in \mathcal{T}$, and a scheduling decision is implemented during each transition $k \in \mathcal{K}$; transition k connects the states observed at stages $t = k$ and $t = k + 1$. A horizon of T transitions, therefore, contains $T + 1$ observed states, and the clock time at stage t is $\theta_t = \theta_0 + t\Delta$.

Figure 1 illustrates the primary decision variable $u_{k,i,j,a}$ on the time–space network, where the horizontal axis is the decision stage and the vertical axis is location. A filled node is the occupied location and an open node is an unselected candidate. During transition k the driver occupies location i, so that $y_{k,i} = 1$, and exactly one transition is selected out of that node: the solid arrow $u_{k,i,j,\mathrm{DR}} = 1$ drives from i to j, giving $y_{k+1,j} = 1$. Driving is the only action that changes location, whereas the non-driving actions stay in place, as with $u_{k+1,j,j,\mathrm{OF}} = 1$. The dashed arrows are the remaining feasible transitions permitted by $q_{k,i,j,a}$ but not chosen.

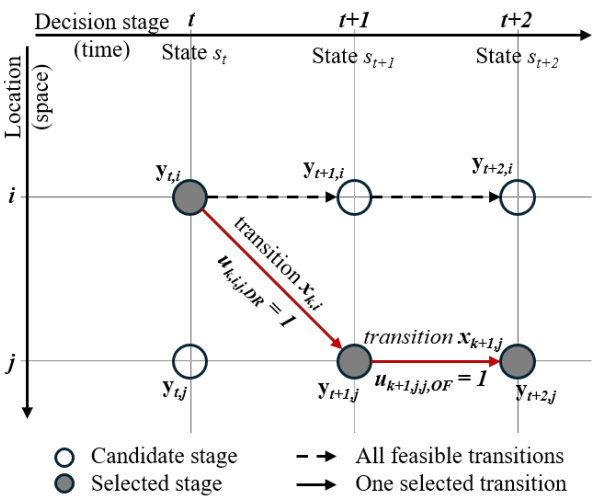


FIGURE 1: Transitions in the time–space network.

At each state-observation stage $t \in \mathcal{T}$, we define a state vector $\mathbf{s}_t$ in Equation (1) that summarizes the information available at the current stage, including the driver's accumulated HOS history, current location, and prevailing operating conditions.

$$\mathbf{s}_t = \Big(\underbrace{R_t^D, R_t^W, R_t^B, R_t^O, w_t}_{\text{HOS history}}, \underbrace{\mathbf{y}_t}_{\text{location}}, \underbrace{P_t, G_t, \Lambda_t, \Omega_t, E_t}_{\text{operating conditions}} \Big) \tag{1}$$

In Equation (1), R_t^τ ($\tau \in \mathcal{H}$) are the four HOS clocks, w_t marks whether the duty window is open, the vector $\mathbf{y}_t = (y_{t,i})_{i\in\mathcal{L}}$ identifies the driver's current location, and $P_t = (\pi_{t,i})_{i\in\mathcal{L}}$, G_t, Λ_t, Ω_t, E_t are the parking availability, traffic conditions, weather conditions, presence of work-zones, and delivery deadlines, respectively. A 30-minute break and a 10-hour rest are not separate actions; they are derived from consecutive non-driving stages. Rerouting is represented by selecting an alternative next location j rather than by a distinct action.

TABLE 1: Notation: sets, parameters, and decision variables.

Symbol	Index/Domain/Value	Definition
Sets and indices		
$\mathcal{L} = \{1, \ldots, L\}$	i, j	Candidate locations; L is the destination.
$\mathcal{T} = \{0, \ldots, T\}$	t	State-observation stages.
$\mathcal{K} = \{0, \ldots, T-1\}$	k	Transition stages; k connects states $t = k$ and $t = k+1$.
$\mathcal{A} = \{\text{DR}, \text{ON}, \text{OF}, \text{SB}\}$	a	Actions: driving, on-duty not driving, off-duty, sleeping.
$\mathcal{H} = \{D, W, B, O\}$	τ	HOS clocks: driving, duty window, break, rolling cycle.
Parameters		
Δ	0.5	Stage duration (h).
N_{10}	20	Stages in a 10-hour rest, $10/\Delta$.
$N_{8\times24}$	384	Stages in the 8-day cycle window, $8(24)/\Delta$.
$q_{k,i,j,a}$	$\{0, 1\}$	Transition feasibility: 1 if $i \to j$ during k under action a.
$p_{t,i}$	$[0, 1]$	Parking availability at location i, stage t.
A_k^{adv}	$\{0, 1\}$	Adverse-condition indicator under 49 CFR 395.1(b)(1).
$\mathbf{R}_k^{\text{lim}}$		HOS limit vector, $(11 + 2A_k^{adv},\ 14 + 2A_k^{adv},\ 8,\ 70)^\top$.
θ_t		Clock time at stage t.
Θ		Delivery deadline.
$\widehat{m}_{k-N_8}^{on}$		On-duty indicator leaving the eight-day window.
$c_{k,i,j,a}$		Generalized transition cost.
λ^L		Lateness penalty weight.
$\boldsymbol{\lambda}^V$		HOS violation penalty weights (one per clock).
M		Sufficiently large constant.
Decision variables		
$u_{k,i,j,a}$	$\{0, 1\}$	Location–action transition selected during k.
$y_{t,i}$	$[0, 1]$	Location occupancy at stage t.
$x_{k,a}$	$[0, 1]$	Duty status selected during k.
$m_k^{dr}, m_k^{on}, m_k^{nd}, m_k^{rs}$	$[0, 1]$	Duty-status indicators for k.
m_k^{win}	$\{0, 1\}$	Duty window active during k.
w_t	$\{0, 1\}$	Duty window open at stage t.
r_t^{10}	$\{0, 1\}$	Qualifying 10-hour rest completed at stage t.
$\mathbf{v}_k$	≥ 0	HOS violation (overrun) slack for k.
b_t	$\{0, 1\}$	Arrival declared at stage t.
ℓ	≥ 0	Delivery lateness.

Pre-Trip Scheduling Model

The pre-trip scheduling model selects the feasible transition sequence of least generalized cost, with the complete formulation provided in Equations (2a) to (2m). The auxiliary relations that define the duty indicators, the qualifying-rest indicator, and the duty-window status are given afterward in Equations (3)–(5), and the clock-change and reset vectors used in (2h) are defined in Equations (6)–(7).

$$\min \quad \sum_{k\in\mathcal{K}}\sum_{i\in\mathcal{L}}\sum_{j\in\mathcal{L}}\sum_{a\in\mathcal{A}} c_{k,i,j,a}\,u_{k,i,j,a} + \lambda^{L}\ell + \sum_{k\in\mathcal{K}} \boldsymbol{\lambda}^{V\top}\mathbf{v}_k \tag{2a}$$

$$\text{s.t.} \quad \textstyle\sum_{i,j,a} u_{k,i,j,a} = 1, \qquad \forall k \tag{2b}$$

$$y_{k,i} = \textstyle\sum_{j,a} u_{k,i,j,a}, \qquad \forall k, i \tag{2c}$$

$$y_{k+1,j} = \textstyle\sum_{i,a} u_{k,i,j,a}, \qquad \forall k, j \tag{2d}$$

$$u_{k,i,j,a} \le q_{k,i,j,a}, \qquad \forall k, i, j, a \tag{2e}$$

$$q_{k,L,j,a} = 0, \qquad \forall k,\ j \neq L,\ a \tag{2f}$$

$$x_{k,a} = \textstyle\sum_{i,j} u_{k,i,j,a}, \qquad \forall k, a \tag{2g}$$

$$\mathbf{R}_{k+1} = \boldsymbol{\Psi}_k\big(\mathbf{R}_k + \Delta\mathbf{d}_k\big),\ \mathbf{R}_k \ge 0 \qquad \forall k \tag{2h}$$

$$\mathbf{R}_k + \Delta\mathbf{g}_k - \mathbf{v}_k \le \mathbf{R}_k^{\lim} + M(1 - m_k^{dr})\mathbf{1}_4,\ \mathbf{v}_k \ge 0 \qquad \forall k \tag{2i}$$

$$x_{k,\mathrm{OF}} + x_{k,\mathrm{SB}} \le \textstyle\sum_i p_{t,i} y_{t,i}, \qquad \forall k \tag{2j}$$

$$\textstyle\sum_t b_t = 1, \quad b_t \le y_{t,L}, \qquad \forall t \tag{2k}$$

$$\ell \ge \textstyle\sum_t \theta_t b_t - \Theta,\ \ell \ge 0, \qquad \forall t \tag{2l}$$

$$u_{k,i,j,a} \in \{0,1\},\ \ b_t \in \{0,1\}, y_{t,i} \in [0,1] \qquad \forall k, i, j, a, t \tag{2m}$$

The objective function in Equation (2a) seeks to minimize the sum of generalized transition cost, lateness, and HOS violation, with $c_{k,i,j,a}$ combining travel time, delay, parking, and operating cost through relative weights. The penalty weights are ordered so that violation dominates lateness, which dominates driving and on-duty cost; only their ordering and ratios are meaningful. In Equations (2c)–(2d) the two sums give the location occupied at the start of k and the location entered at $k+1$. In Equation (2e) for the feasibility parameter, $q_{k,i,i,\mathrm{DR}}$ representing 30 minutes of driving that does not complete a segment and $q_{k,i,j,a} = 0$ for $i \neq j$ when $a \neq \mathrm{DR}$. The remaining constraints make the destination absorbing (2f), extract the duty status (2g), enforce parking availability for off-duty and sleeper-berth stages (2j), declare a single arrival at the destination (2k), and measure lateness against the deadline Θ (2l). The clock update (2h) and soft limit (2i) are described in the following subsections.

The duty status selected in Equation (2g) is mapped to the stage-level indicators that drive the HOS clocks, as defined in Equation (3). Driving is the only status advancing the driving and break clocks, on-duty time advances the rolling cycle, and off-duty and sleeper-berth time contribute to a qualifying rest. Because each stage lasts $\Delta = 0.5$ h, a single non-driving stage ($m_k^{nd} = 1$) is one complete 0.5 h period and resets the break clock directly. The sleeper-berth split provision is not modeled, so off-duty and sleeper-berth time enter identically.

$$\begin{cases} m_k^{dr} = x_{k,\mathrm{DR}}, \\ m_k^{on} = x_{k,\mathrm{DR}} + x_{k,\mathrm{ON}}, \\ m_k^{nd} = x_{k,\mathrm{ON}} + x_{k,\mathrm{OF}} + x_{k,\mathrm{SB}}, \\ m_k^{rs} = x_{k,\mathrm{OF}} + x_{k,\mathrm{SB}}, \end{cases} \qquad \forall k \in \mathcal{K} \tag{3}$$

A qualifying 10-hour rest completes after $N_{10} = 20$ consecutive stages with $m^{rs} = 1$, as expressed by the rest-completion indicator in Equation (4). This is linearized by $r_t^{10} \le m_k^{rs}$ for each

$k \in \{t - N_{10}, \ldots, t-1\}$ and $r_t^{10} \geq \sum_k m_k^{rs} - (N_{10} - 1)$.

$$r_t^{10} = \mathbb{I}\Big\{ \textstyle\sum_{t-N_{10}}^{t-1} m_k^{rs} = N_{10} \Big\}, \quad \forall t \tag{4}$$

The 14-hour window opens when the driver first goes on duty after a qualifying rest and runs continuously until the next one. It is carried in w_k and updated with the window indicator by Equation (5).

$$\begin{cases} m_k^{win} = \max\{w_k,\, m_k^{on}\}, \\ w_{k+1} = (1 - r_{k+1}^{10})\, m_k^{win}, \end{cases} \qquad \forall k \in \mathcal{K} \tag{5}$$

The four clocks are collected in $\mathbf{R}_t = (R_t^D, R_t^W, R_t^B, R_t^O)^\top$, initialized from ELD history. Their update (2h) uses two vectors: the net clock-change vector in Equation (6) and the reset matrix in Equation (7).

$$\mathbf{d}_k = \left(m_k^{dr},\, m_k^{win},\, m_k^{dr},\, m_k^{on} - \widehat{m}_{k-N_8}^{on}\right)^\top \tag{6}$$

$$\mathbf{\Psi}_k = \operatorname{diag}\left(1 - r_{k+1}^{10},\, 1 - r_{k+1}^{10},\, 1 - m_k^{nd},\, 1\right) \tag{7}$$

The fourth component of $\mathbf{d}_k$ advances the rolling cycle by adding the current on-duty stage and removing the stage leaving the 8-day window. Since every trip here satisfies $T \leq N_8$, the expiring term $\widehat{m}_{k-N_8}^{on}$ is fixed by pre-horizon history and is exogenous, which is what lets the rolling cycle be represented without carrying eight days of duty history in the state. In $\mathbf{\Psi}_k$, a qualifying rest resets the driving and window clocks, a non-driving stage resets the break clock, and the rolling cycle is never reset by these short-term conditions. Applying (2h) adds the net change and then the resets; its fourth row is equivalent to $R_{k+1}^O = R_k^O + \Delta(m_k^{on} - \widehat{m}_{k-N_8}^{on})$.

Driving eligibility is checked by the soft limit (2i), which uses the HOS-use vector $\mathbf{g}_k = (m_k^{dr}, m_k^{win}, m_k^{dr}, m_k^{on})^\top$ and the limit vector $\mathbf{R}_k^{\lim} = (11 + 2A_k^{adv},\, 14 + 2A_k^{adv},\, 8,\, 70)^\top$. The vector $\mathbf{g}_k$ differs from $\mathbf{d}_k$ only in the fourth component, because eligibility depends on the on-duty time a transition adds, not on time expiring from the window; the adverse indicator A_k^{adv} extends the driving and window limits by two hours. The nonnegative slack $\mathbf{v}_k$ absorbs any overrun, so no driving transition is infeasible under the regulation; each hour of overrun is instead charged at $\boldsymbol{\lambda}^V$ through (2a). The big-M term relaxes the check for non-driving actions ($m_k^{dr} = 0$), and $\boldsymbol{\lambda}^V \to \infty$ forces $\mathbf{v}_k = 0$, recovering the hard-constraint case. The joint effect of these three limits, the 11-hour driving limit, the 14-hour window, and the 8-hour break, on whether a trip can be completed within the current shift is shown in Figure 2.

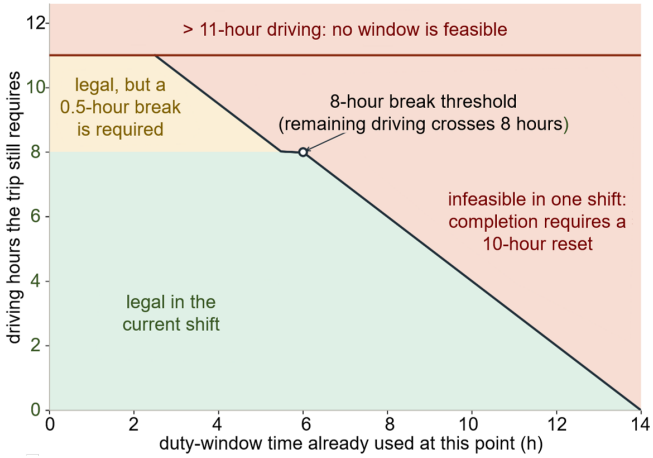


FIGURE 2: HOS-implied feasibility regions.

The products of the two decision variables are linearized with standard big-M and indicator constraints, so the model stays a mixed-integer linear program. The formulation uses the 70-hour, 8-day cycle limit, which spans $N_{8\times24} = 384$ stages. The 34-hour restart and sleeper-berth split provision are not modeled and are addressed in the limitations. Solving Equations (2) yields the pre-trip schedule $\mathcal{P}^0 = \{(j_k^0, a_k^0) : k \in \mathcal{K}\}$.

In-Trip Scheduling Model

The pre-trip plan is determined from forecasts. During the trip, the in-trip scheduling model revises the remaining schedule as the observed state $\mathbf{s}_t$ becomes available. It is a finite-horizon stochastic dynamic program, stated in Equations (8) and solved by exact backward induction.

$$C_t(\mathbf{s}) = \min_{(j,a)} \widetilde{c}_k(\mathbf{s}, j, a) + \lambda^{dev}\mathbb{I}\big[(j,a) \neq (j_k^0, a_k^0)\big] + \mathbb{E}_{\xi_k}\big[C_{t+1}(f_k(\mathbf{s}, j, a, \xi_k))\big] \tag{8a}$$

$$\text{s.t.} \quad (j,a) \in \mathcal{D}_k(\mathbf{s}_t; \alpha^{HOS}) \tag{8b}$$

$$\mathcal{D}_k(\mathbf{s}_t; \alpha^{HOS}) = \big\{(j,a) \in \mathcal{D}_k(\mathbf{s}_t) : p_k^{viol}(j,a \mid \mathbf{s}_t) \leq \alpha^{HOS}\big\} \tag{8c}$$

$$p_k^{viol}(j,a \mid \mathbf{s}_t) = \Pr_{\xi_k}\big[f_k(\mathbf{s}_t, j, a, \xi_k) \notin \mathcal{S}^{HOS}\big] \tag{8d}$$

$$C_T(\mathbf{s}) = \Phi(\mathbf{s}) \tag{8e}$$

The three terms in (8a) are the one-step transition cost, including any HOS violation penalty; a penalty λ^{dev} on deviating from the pre-trip schedule, so the plan is retained unless revision is worthwhile; and the expected cost-to-go, where f_k maps the state and decision to the next state under uncertainty ξ_k, taken here as realized delay and parking availability. Because the HOS constraint is soft, an action that would overrun a clock remains admissible, and the overrun enters $\widetilde{c}_k$ rather than removing the action. The feasible set $\mathcal{D}_k(\mathbf{s}_t)$ applies the same transition, HOS, parking, and destination conditions as the base model at the observed state. The risk filter (8c) then retains only transitions whose one-step violation probability (8d) is within the tolerance α^{HOS}, and setting $\alpha^{HOS} = 1$ leaves the feasible set unscreened, since $\mathcal{D}_k(\mathbf{s}_t; 1) = \mathcal{D}_k(\mathbf{s}_t)$. Because R^O is exogenous, the state carried in the recursion reduces to $(R_t^D, R_t^W, R_t^B, w_t, \mathbf{y}_t)$, which permits an exact solution rather than an approximation architecture.

The remaining scheduling flexibility at state k is summarized by the fraction of feasible transitions that remain after screening, i.e., $\big|\mathcal{D}_k(\mathbf{s}_k; \alpha^{\text{HOS}})\big| \,/\, |\mathcal{D}_k(\mathbf{s}_k)|$, which is a ratio of set sizes that indicates how constrained the current state has become. This measure differs from p_k^{viol}, which evaluates the violation risk associated with a single candidate transition. A state may contain one high-risk transition while still retaining many acceptable alternatives.

Solution Algorithm and Performance Measures

The pre-trip model (2) is a mixed-integer linear program and is solved to optimality with the COIN-OR Branch and Cut solver through the PuLP interface. The adaptive model (8) is a Bellman recursion and is solved by backward induction over the reduced state $(R_t^D, R_t^W, R_t^B, w_t, \mathbf{y}_t)$ and evaluating the expectation over the finite support of ξ_k at each stage. Because the support is discrete and the state space is enumerable, no approximation or sampling is required. HOS violation is reported two ways: the violation rate, the fraction of scenarios with any overrun at any stage, and the violation hours, the mean total overrun per scenario. On-time completion is the fraction of scenarios arriving by the deadline. Reporting rate with hours separates how often a schedule violates from how severely, and pairing either with on-time completion exposes the compliance–timeliness trade-off that the soft formulation represents.

DATA AND EXPERIMENTAL DESIGN

Data Sources

This study draws on three data sources. A synthesized corridor instance provides full experimental control for characterizing the base scheduling model and its sensitivity. ELD data from a trucking operation provide real-world duty-status sequences and are used to reconstruct driver schedules, HOS clocks, driving blocks, break timing, rest opportunities, and multi-day work-rest patterns. Simulator experiment data collected from CMV drivers provide controlled observations of driver decision-making under designed operating conditions, including short-break decisions, drop-off or delivery-related events, adverse weather, work zones, and other disruptions. The three sources play complementary rather than overlapping roles, as summarized in Table 2.

TABLE 2: Supporting model components from synthesized, ELD log, and simulator test data.

Model component	Synthesized data	ELD log data	Simulator test data
R^B 30-min break	Yes	Yes	Yes
R^D 11-h driving	Yes	Yes (max = 10.37 h)	Partial
R^W 14-h window	Yes	Yes	Partial
R^O 70-h/8-day cycle	No	Yes (peaks at 72.04 h)	No
10-h rest	Yes	Yes (29 observed)	No
34-h restart	No	Partial (1 observed for 38.95 h)	No
Split sleeper	No	No	No
Parking $\pi_{k,i}$	Yes (20%–80%)	Yes (110 real facility choices)	No
Vehicle dynamics	No	No	Yes (10 Hz)
Controlled manipulation	Yes	No	Yes (S1/S2)
Counterfactual DPs	Yes	No	Yes

Synthesized Data and Parameter Setup

The base and adaptive scheduling models are implemented on a synthesized dataset to isolate how HOS limits, delivery timing, parking availability, and weather jointly shape schedule feasibility. A full factorial design varies four operational factors and, for each factor combination, draws random realizations of parking and weather to measure how often a schedule fails. The swept factors and their levels are trip length (i.e., 7 to 15 driving hours), delivery-window slack over the minimum feasible time (i.e., 0 to 3 hours), the fraction of usable parking areas (i.e., 20% to 80%), and adverse-weather fraction (i.e., 0% to 60% of the trip). The full design contains 1,512 factor cells, each evaluated over 150 stochastic scenarios, for roughly 2.3×10^5 simulated trips. A trip is recorded as a failure when the driver is forced into an HOS violation or cannot complete the trip legally within the delivery window.

ELD Data Description

The ELD dataset comprises 553 duty-status events recorded over 30 consecutive days from a truck unit operating a cross-country route through Nebraska, Iowa, Wyoming, Oregon, and California. Total driving time is 227.6 hours across 110 unique stop locations. Figure 3(a) shows the ELD duty-status record.

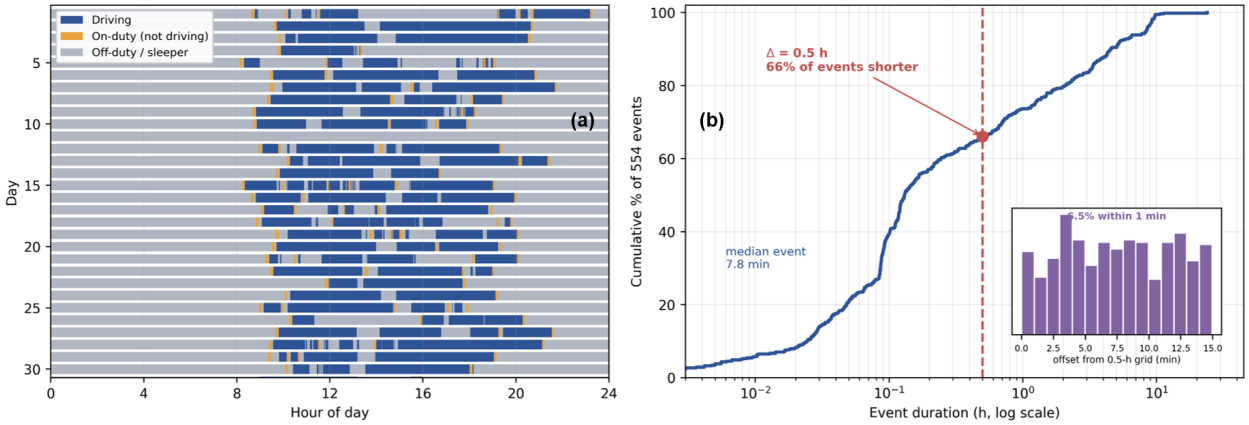


FIGURE 3: ELD data records with (a) duty-status structure over the 30-day horizon and (b) event granularity relative to the 30-minute stage resolution.

Long rests in this ELD record are logged as off-duty rather than sleeper berth. Qualifying rest is therefore identified by duration rather than by label, and the reconstructed clocks depend on how rest is coded. Under a strict rest-coding rule (at least 10 hours reset the driving and duty-window clocks), 15 periods violate the 14-hour window. Under a relaxed rest-coding rule (at least 8 hours as qualifying rests), only 2 periods violate the 14-hour window. Therefore, 13 apparent violations are an artifact of duty-status coding rather than of driving behavior. The rolling 8-day on-duty time peaks at 72.0 hours and exceeds the 70-hour cycle limit on 3 days, and 8 driving periods exceed the 8-hour limit since the last qualifying break, reaching 8.76 hours. The 11-hour driving limit is never violated, with a maximum of 10.37 hours. As can be seen in Figure 3(b), most events are shorter than a single 30-minute stage and status changes do not align to the clock grid, which motivates the stage-based state reconstruction used in the model.

Simulator Experiment Data

Twelve CMV drivers each completed two scenarios in a driving simulator in random order. The two scenarios S1 and S2 traverse the same roadway, as shown in Figure 4. The scenarios differ in speed regime rather than distance: work zones and rain hold S2 drivers to a lower speed, so over the shared segment they accumulate roughly 7.6 to 8.0 driving hours against 5.2 in S1 and reach the same geographic point against the eight-hour break threshold. Rest facilities are placed at designated decision points, where the driver either stops to rest or proceeds. This binary choice is the primary behavioral outcome. Each run began with an assumed five hours already on the driving and duty-window clocks. Telemetry is logged at approximately 10 Hz and records vehicle pose, speed, ignition state, and duty status. Time is compressed by design, one simulator second representing one real-world minute, so a run of 9 to 15 wall-clock minutes corresponds to a duty day of 5.4 to 14.6 real hours, placing each run on the hours-of-service scale the model uses.

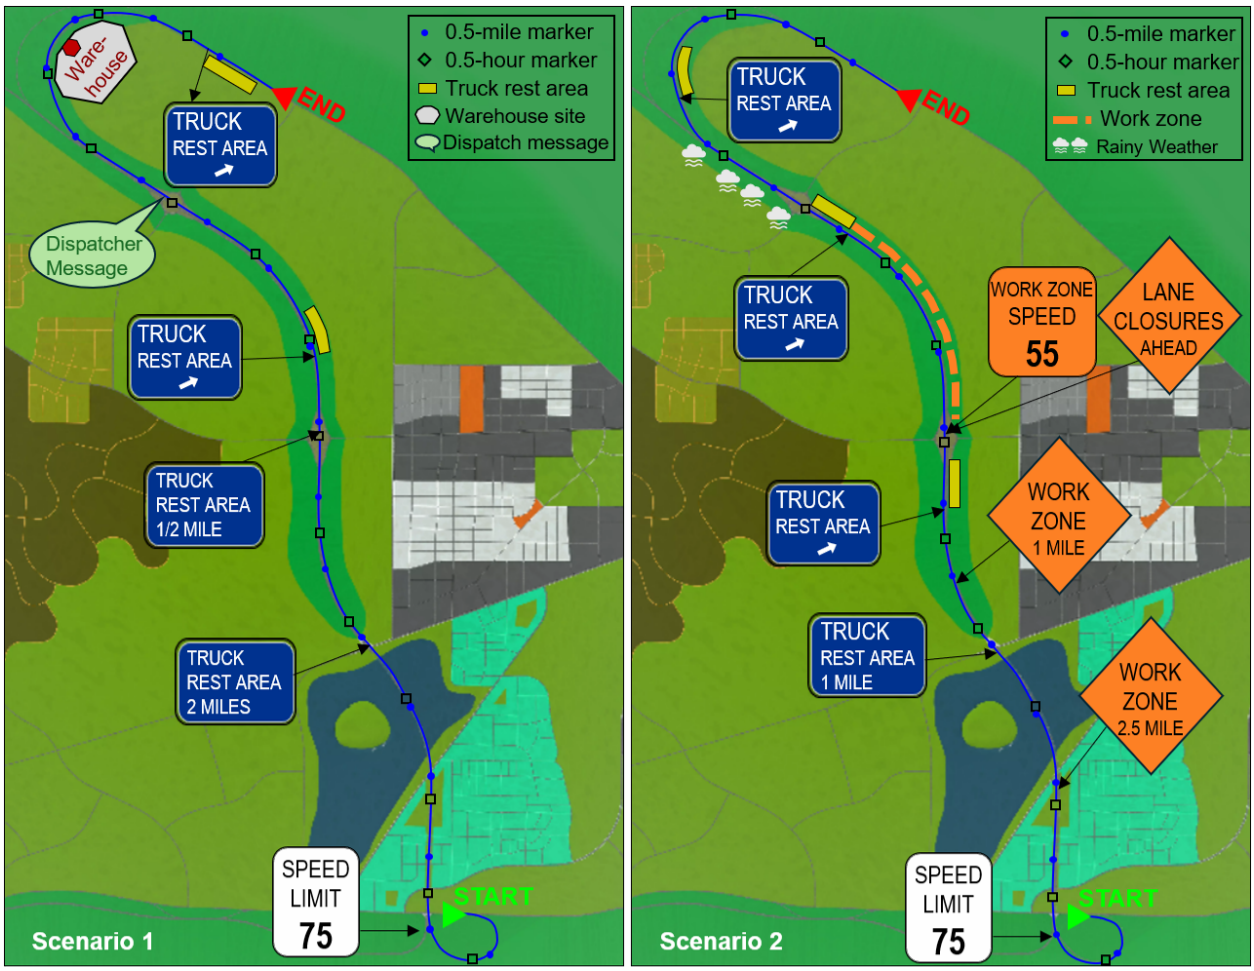


FIGURE 4: Two scenario settings in the simulator.

The two data sources instantiate the model's uncertainty term ξ_k in complementary ways. The ELD record supports state reconstruction but contains no controlled variation, so ξ_k cannot be observed from it directly; the simulator supplies ξ_k through controlled events, including reduced speed under weather, work-zone delay, unavailable parking at a planned stop, and a short-break need. These events enter the model through the transition cost $c_{k,i,j,a}$, the feasibility parameter $q_{k,i,j,a}$, the parking availability $p_{t,i}$, or the observed state $\mathbf{s}_t$, and the adaptive model then determines whether the driver should continue driving or take a rest.

RESULTS

Simulation-based Model Output and Sensitivity

Table 3 reports the base MIP for a synthesized corridor instance at five trip lengths, with 0.50 parking availability, no adverse weather, a delivery deadline of 1.6× driving hours, and an assumed speed of 75 mph. Every instance returns a legal plan that arrives within the deadline, so the violation slack stays inactive and no overrun is incurred. The objective rises almost linearly with trip length, from 14.0 at 7 driving hours to 30.1 at 15 driving hours, because driving cost accumulates one unit per half-hour stage while rest adds only a small idle cost. Trips of 9 driving hours and longer insert a single 30-minute break to satisfy the 8-hour break limit, and arrival time tracks the driving

requirement plus that break, remaining inside the deadline in every case.

The absence of any violation is a property of the cost ordering rather than of a loose instance. Because the violation penalty dominates lateness, which in turn dominates driving and idle cost, the model accepts a violation only when no legal plan exists. When the deadline is tightened to the driving time itself, the model responds by admitting a small lateness rather than an overrun, since the lateness penalty is the cheaper of the two. The base MIP therefore behaves as a compliance-first planner: it produces a legal pre-trip schedule whenever one is available and reserves violation for cases that the schedule alone cannot resolve. These are handled by the adaptive model, where realized uncertainty during the trip can force an overrun that no pre-trip plan anticipated.

TABLE 3: Base MIP solved to optimality for a synthesized corridor instance.

Trip length (h)	Objective cost*	Driving (h)	Rest (h)	Violation (h)	Lateness (h)	Arrival (h)
7	14.0	7.0	0.0	0.0	0.0	11.0
9	18.1	9.0	0.5	0.0	0.0	13.5
11	22.1	11.0	0.5	0.0	0.0	17.5
13	26.1	13.0	0.5	0.0	0.0	20.5
15	30.1	15.0	0.5	0.0	0.0	24.0

*Objective cost is dimensionless, and only the ordering and ratios of the weights are meaningful.

The base MIP produces a legal pre-trip plan, but that plan is computed from forecast conditions and is executed open-loop. When delay and parking are realized during the trip, a plan with little slack is driven to the edge of the HOS limits, and the adaptive model revises the remaining sequence from the observed state. Table 4 evaluates this revision by comparing the adaptive scheduling against the open-loop baseline scheduling under identical realized uncertainty. The comparisons sweep four factors: trip length, parking availability, per-segment delay probability, and the risk threshold α^{HOS}, one factor at a time. Each factor varies around a reference operating point while the others are held fixed. Across every condition the adaptive model holds the HOS violation rate at or below 3.5%, against a baseline that exceeds 90% once the trip approaches the driving limit. The two outcome columns must be read together, because the baseline sustains its high on-time rate by continuing to drive past the limits rather than stopping to rest, so its near-complete on-time record is obtained through violations. The adaptive model instead stops when required, which is why its violation rate is low and its on-time rate is the quantity that varies. The gain in legality is therefore robust to operating conditions; what changes with those conditions is its cost in timeliness.

That cost tracks how tightly each factor constrains the schedule. When a trip fits within a single shift, on-time completion stays above 89% through 9 driving hours, so violations are removed at little delay. When the trip is longer, a legal schedule must insert a 10-hour reset that pushes arrival past the deadline, and on-time completion falls sharply, to 41.5% at 10 driving hours and to near zero beyond. The same pattern holds for the other two conditions: on-time completion falls as parking grows scarce, from 42.5% at 50% availability to 29.5% at 20%, and as delays become frequent, from 92.0% at a delay probability of 0.05 to 12.5% at 0.25 and to zero above it. In each case the adaptive model converts a near-certain violation into a late but legal arrival, and the residual lateness is largest exactly where the trip is least recoverable by scheduling alone.

The risk threshold α^{HOS} behaves as a single transition rather than a continuous dial. Below roughly 0.13 the screen is strict enough to drive the violation rate to zero, but it forces earlier rests

and holds on-time completion to 20.5%; above that point a set of borderline transitions becomes admissible at once, and the outcome settles at a 3.5% violation rate with 41.5% on-time completion, unchanged across $\alpha^{HOS} \in \{0.2, 0.3, 0.4, 0.5\}$. The flat response on this plateau means the reported results do not depend on the precise value of α^{HOS} near the default of 0.3. The sharpness of the single transition reflects the discrete delay and parking outcomes used here; under richer uncertainty the response would smooth, though the trade-off between legality and timeliness would remain.

TABLE 4: HOS violation and on-time completion under baseline plan and adaptive scheduling

	HOS violation (%)		On-time (%)	
Factor	Baseline	Adaptive	Baseline	Adaptive
Trip length (driving hours)				
6 h	2.5	0.0	100.0	98.5
7 h	36.5	0.0	100.0	98.0
8 h	91.0	0.5	100.0	98.0
9 h	93.0	3.0	100.0	89.0
10 h	93.0	3.5	100.0	41.5
11 h	96.5	0.5	100.0	3.5
12 h	93.3	0.0	93.3	0.0
Parking availability				
20%	98.5	2.5	100.0	29.5
50%	95.0	3.0	100.0	42.5
80%	92.5	3.0	100.0	41.0
Per-segment delay probability				
0.05	72.5	1.5	100.0	92.0
0.25	99.0	1.5	99.5	12.5
0.45	100.0	0.0	88.0	0.0
0.65	100.0	0.0	24.0	0.0
Risk threshold α^{HOS}				
0.1	93.0	0.0	100.0	20.5
0.2	93.0	3.5	100.0	41.5
0.3	93.0	3.5	100.0	41.5
0.4	93.0	3.5	100.0	41.5
0.5	93.0	3.5	100.0	41.5

In summary, the results delimit where adaptive scheduling is decisive. On trips that remain feasible within a shift it removes violations at little cost to on-time performance, which is the operationally valuable regime. On trips whose driving requirement, parking scarcity, or delay exposure exceed what a single shift can absorb, it still guarantees legality, but only by trading timeliness for compliance, and the residual lateness signals a dispatch-level problem that scheduling alone cannot resolve.

ELD-based State Reconstruction and HOS Exposure

All results from ELD data are reconstructed from the raw log using the clock-update logic defined in the Methods section. Table 5 reports the reconstructed HOS violation checks under two rest-coding rules. Strict coding uses a restrictive interpretation of which non-driving periods qualify as rest,

while relaxed coding treats any continuous non-driving period as rest once the duration requirement is met. The reconstructed driving clock is stable under both rules: no driving-limit violation is detected, and the maximum reconstructed value is 10.37 hours, safely below the 11-hour limit. The drive-since-break clock is likewise identical under both rules, with the same eight violations and a maximum of 8.76 hours.

The duty-window clock is sensitive to the rest-coding rule. Under strict coding it shows 15 violations with a maximum of 23.27 hours; under relaxed coding the count falls to 2 with a maximum of 14.65 hours. Without crediting the long off-duty periods as qualifying rest, the window does not reset and accumulates across activity that is in fact compliant. The driving clock is unaffected by the rule, so the coding choice changes the apparent compliance record through the duty window rather than through driving. The driving clock, unlike the duty window, is not the binding constraint. Mean daily driving is 7.6 hours and reaches or exceeds 8 hours on 16 of the 30 days, with a single-day maximum of 11.2 hours, but the 11-hour limit is not exceeded. The binding pressure is instead the multi-day accumulation of on-duty time: the 8-day rolling total reaches 72.04 hours and exceeds the 70-hour limit on three days. This is why the rolling-cycle clock R_k^O is carried in the state vector rather than modeling each day in isolation.

TABLE 5: HOS violation checks under strict and relaxed rest coding

Clock	Limit (h)	Strict coding	Relaxed coding
Drive since break	8	8 violations, max = 8.76	8 violations, max = 8.76
Driving	11	0 violations, max = 10.37	0 violations, max = 10.37
Duty window	14	15 violations, max = 23.27	2 violations, max = 14.65
Cycle, 8 day	70	3 days over, max = 72.04	3 days over, max = 72.04

These results show that the raw ELD record can be converted into the state sequence the base model requires: the activity log yields observed actions and candidate stopping locations, the duration-based rest logic yields the qualifying-break and qualifying-rest indicators, and the clock-update logic yields the four HOS clocks over time. The reconstruction also localizes the periods of limited HOS margin, those approaching the 8-hour break, 14-hour window, and 70-hour cycle limits, where rescheduling has the most value.

The ELD records support reconstruction of HOS states and identification of scheduling pressure, but they contain no controlled variation in weather, work-zone, or delivery conditions and so cannot identify how drivers trade off stopping against proceeding. The ELD data are therefore used for state reconstruction and HOS-exposure analysis, while stop-versus-proceed behavior in the adaptive model is evaluated separately under controlled conditions.

Simulator Driving Test Results

Twelve drivers each drove the two routes in the simulator, wherein one simulated minute represents one hour of driving, and each driver starts with five hours already on the driving and duty-window clocks. The routes differ in rest-area structure. Scenario 1 offers a single mid-route rest area, while Scenario 2 offers three. At each rest area, the driver either stops or continues. Figure 5 compares the stop-or-go recommendation of the pre-trip MIP and the in-trip DP against the observed decision. For each data point, the color marks the observed decision with red for a stop and green for a go. The data also indicate the driving hours versus duty-window hours (or their ratios %) when the stop-or-go decision was made.

Scenario 1 provides one rest area near the end of a long drive, so the drivers face a single late stop-or-continue decision. Eleven of the twelve drivers produced a complete route. At that rest area, 6 of the 11 stopped, and the decision follows the driving clock only weakly. Drivers who stopped had used 10.2 h of driving on average against 9.4 h for those who continued, but the two groups overlap. With a single decision point and no earlier rest options, the accumulated clock alone does not separate the two choices.

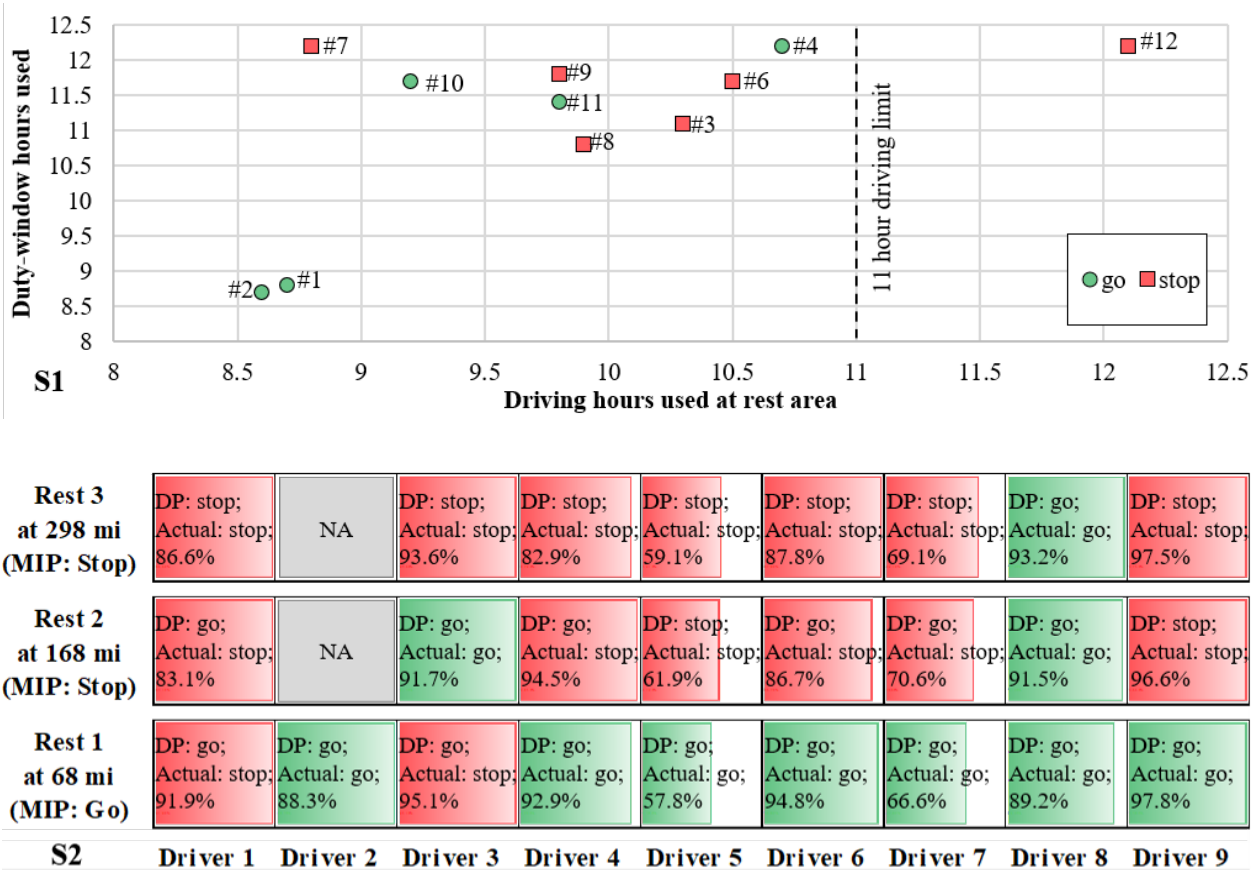


FIGURE 5: Model recommendation versus observed driver decision at each rest area. Upper: Scenario 1, Lower: Scenario 2. Data from Drivers #2 and #5 are incomplete.

Scenario 2 offers three rest areas, and here the models and the drivers agree at most points. The MIP recommendation matches the observed decision at 20 of 25 points and the DP at 19 of 25. Both reproduce the main pattern. Drivers continue at the first rest area, where 2 of 9 stop, and stop at the later two, where 6 of 8 stop at the second and 7 of 8 at the third rest area. The disagreements are one-directional. Wherever a model differs from the driver, the model recommends going while the driver stops, and this occurs when the driving clock still has one to two hours of margin. Drivers stop earlier, which leaves a rest buffer that the cost objective does not reward.

The driving-to-window ratio describes a driver trait, not the stop decision. It is stable within a driver across the trip and ranges from about 58% to 97%, measuring how much of the elapsed window each driver spends driving. It does not predict stopping. A driver near 90% continues at every rest area while a driver near 60% stops at the later two rest areas. A high ratio means the driving and window limits are reached together, and a low ratio means the window is reached first.

The contrast between the two scenarios is the useful result. The models track observed behavior when the route offers several rest areas, and stopping is driven by accumulating HOS pressure. When the route concentrates the choice into one late decision, as in Scenario 1, behavior is governed more by individual preference than by the clock, which the aggregate model does not attempt to capture. The simulator evidence for the scheduling models therefore rests on Scenario 2, with Scenario 1 marking the condition under which route structure rather than HOS state dominates the stopping decision.

CONCLUSIONS

We developed a two-part optimization framework for HOS-constrained scheduling of CMV drivers. A pre-trip model plans a legal schedule before the trip, and an in-trip model revises the remaining trip from the observed state as conditions change. We represent HOS limits as soft constraints, so a violation is penalized rather than forbidden, which lets us report both whether and by how much a schedule breaks the regulation. We solve the pre-trip MIP model to the best possible schedule and the in-trip DP model by working backward from the end of the trip over a reduced set of states.

Our results support three findings. First, on synthesized trips, the adaptive model holds the HOS violation rate at or below 3.5% across trip length, parking availability, and delay probability, against a baseline that exceeds 90% once the trip approaches the driving limit. On short trips that fit one shift, the adaptive model removes violations with little effect on on-time arrival, while on long trips that do not, it turns a near-certain violation into a late but legal arrival, which is where it helps most. Second, a raw ELD record can be reconstructed into the HOS clock states the model requires, and the reconstruction shows that the limit a real duty record reaches first is the timing of rest and the buildup of on-duty time over several days, not the daily driving limit. Third, the driving-simulator study adds a behavioral check. The model recommendations match observed stop-or-go decisions at most rest areas when the route has several rest areas, and drivers tend to stop earlier than the cost-optimal plan requires.

A few limitations should be noted. First, the optimization runs on a single-corridor network with locations in trip order, so a general road network would require the full location-to-location transition variables. Second, the reported magnitudes are properties of the synthesized instances and the assumed relative cost weights. The ordering and shape of the results are informative, but the specific percentages are not calibrated to a particular carrier. Third, the adaptive dynamic program is solved on a reduced single-chain state that captures the HOS clocks but not the full multi-location network of the base MIP. We leave addressing these considerations to future work.

The framework has distinct value for several stakeholders. For carriers and dispatchers, the results indicate that HOS compliance is improved most by pairing conservative trip assignment with reactive rescheduling rather than by either alone, because the adaptive model recovers marginally feasible trips but cannot rescue trips whose driving requirement exceeds a single shift. For drivers, the adaptive policy provides a decision at each stage, whether to continue, break, rest, or wait, that keeps the trip legal without depending on a plan fixed before departure. For regulators and safety analysts, the reconstruction shows that duty-status coding conventions can change the apparent compliance record without any change in driving behavior, which is relevant to how ELD records are audited. For researchers, the soft-constraint formulation and the exact adaptive recursion provide a structure that can be extended to a general network, to richer uncertainty, and to larger driver samples as data become available.

ACKNOWLEDGMENTS

This material is based upon work supported by the Federal Motor Carrier Safety Administration under grant FM-MHP-0889. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the view of the Federal Motor Carrier Safety Administration and/or the U.S. Department of Transportation.

AUTHOR CONTRIBUTIONS

The authors confirm contributions to the paper as follows: study conception and design: L Zhao, J Hawkins, N Huynh; data collection: H Afzal, J Cao, L Zhao; analysis and interpretation of results: L Zhao, H Afzal, J Hawkins, and N Huynh; draft manuscript preparation: L Zhao, H Afzal, and J Cao; revise manuscript: All authors. All authors reviewed the results and approved the final version of the manuscript.

DECLARATION OF CONFLICTING INTERESTS

The authors declare no conflicts of interest with any other entities or researchers.

REFERENCES

FMCSA, *Summary of Hours of Service Regulations*. U.S. Department of Transportation, Washington, DC, 2022, accessed July 21, 2026.

Soccolich, S. A., M. Blanco, R. J. Hanowski, R. L. Olson, J. F. Morgan, F. Guo, and S.-C. Wu, An Analysis of Driving and Working Hour on Commercial Motor Vehicle Driver Safety Using Naturalistic Data Collection. *Accident Analysis & Prevention*, Vol. 58, 2013, pp. 249–258.

Goel, A. and S. Irnich, An Exact Method for Vehicle Routing and Truck Driver Scheduling Problems. *Transportation Science*, Vol. 51, No. 2, 2017, pp. 737–754.

Pilz, D., S. Schwerdfeger, and N. Boysen, Make or Break: Coordinated Assignment of Parking Space for Breaks and Rest Periods in Long-Haul Trucking. *Transportation Research Part B: Methodological*, Vol. 164, 2022, pp. 45–64.

Pillac, V., M. Gendreau, C. Guéret, and A. L. Medaglia, A Review of Dynamic Vehicle Routing Problems. *European Journal of Operational Research*, Vol. 225, No. 1, 2013, pp. 1–11.

Goel, A. and L. Kok, Truck Driver Scheduling in the United States. *Transportation Science*, Vol. 46, No. 3, 2012, pp. 317–326.

Blanco, M., R. J. Hanowski, R. L. Olson, J. F. Morgan, S. A. Soccolich, S.-C. Wu, and F. Guo, *The Impact of Driving, Non-Driving Work, and Rest Breaks on Driving Performance in Commercial Motor Vehicle Operations*. Federal Motor Carrier Safety Administration, U.S. Department of Transportation, Washington, DC, 2011.

Jovanis, P. P., K.-F. Wu, and C. Chen, Effects of Hours of Service and Driving Patterns on Motor Carrier Crashes. *Transportation Research Record: Journal of the Transportation Research Board*, Vol. 2281, No. 1, 2012, pp. 119–127.

Chen, C. and Y. Xie, The Impacts of Multiple Rest-Break Periods on Commercial Truck Driver's Crash Risk. *Journal of Safety Research*, Vol. 48, 2014a, pp. 87–93.

Chen, C. and Y. Xie, Modeling the Safety Impacts of Driving Hours and Rest Breaks on Truck Drivers Considering Time-Dependent Covariates. *Journal of Safety Research*, Vol. 51, 2014b, pp. 57–63.

Goel, A., Truck Driver Scheduling in the European Union. *Transportation Science*, Vol. 44, No. 4, 2010, pp. 429–441.

1 Rancourt, M.-È., J.-F. Cordeau, and G. Laporte, Long-Haul Vehicle Routing and Scheduling with
2 Working Hour Rules. *Transportation Science*, Vol. 47, No. 1, 2013, pp. 81–107.
3 Goel, A. and T. Vidal, Hours of Service Regulations in Road Freight Transport: An Optimization-
4 Based International Assessment. *Transportation Science*, Vol. 48, No. 3, 2014, pp. 391–412.
5 Kok, A. L., C. M. Meyer, H. Kopfer, and J. M. J. Schutten, A Dynamic Programming Heuristic
6 for the Vehicle Routing Problem with Time Windows and European Community Social
7 Legislation. *Transportation Science*, Vol. 44, No. 4, 2010, pp. 442–454.
8 Kok, A. L., E. W. Hans, and J. M. J. Schutten, Vehicle Routing under Time-Dependent Travel
9 Times: The Impact of Congestion Avoidance. *Computers & Operations Research*, Vol. 39,
10 No. 5, 2012, pp. 910–918.
11 Vital, F. and P. Ioannou, Scheduling and Shortest Path for Trucks with Working Hours and Parking
12 Availability Constraints. *Transportation Research Part B: Methodological*, Vol. 148, 2021,
13 pp. 1–37.
14 Federal Highway Administration, *Jason's Law Commercial Motor Vehicle Parking Survey and*
15 *Comparative Assessment*. U.S. Department of Transportation, Federal Highway Adminis-
16 tration, Washington, DC, 2020.
17 Vital, F. and P. A. Ioannou, Long-Haul Truck Scheduling with Driving Hours and Parking Avail-
18 ability Constraints. In *2019 IEEE Intelligent Vehicles Symposium*, IEEE, 2019, pp. 620–625.
19 Boris, C. and R. Brewster, A Comparative Analysis of Truck Parking Travel Diary Data. *Trans-*
20 *portation Research Record*, Vol. 2672, No. 9, 2018, pp. 242–248.
21 Gaddy, A. C., S. Hernandez, and S. G. Nurre, Incorporating Truck Parking and Hours of Service
22 (HOS) into a Truck Routing Heuristic. In *Proceedings of the 97th Annual Meeting of the*
23 *Transportation Research Board*, Washington, DC, 2018, 18-00852.
24 Vital, F., P. A. Ioannou, and A. Gupta, Survey on Intelligent Truck Parking: Issues and Approaches.
25 *IEEE Intelligent Transportation Systems Magazine*, Vol. 13, No. 4, 2021, pp. 31–44.
26 Morris, T. L., V. Morellas, N. P. Papanikolopoulos, and M. R. Gallagher, A Region-Wide Deploy-
27 ment of an Automated Real-Time Truck Parking Availability Detection and Information
28 Dissemination System. *IFAC-PapersOnLine*, Vol. 58, No. 10, 2024, pp. 267–272.
29 Taş, D., N. Dellaert, T. van Woensel, and T. de Kok, Vehicle Routing Problem with Stochastic
30 Travel Times Including Soft Time Windows and Service Costs. *Computers & Operations*
31 *Research*, Vol. 40, No. 1, 2013, pp. 214–224.
32 Taş, D., N. Dellaert, T. van Woensel, and T. de Kok, The Time-Dependent Vehicle Routing Problem
33 with Soft Time Windows and Stochastic Travel Times. *Transportation Research Part C:*
34 *Emerging Technologies*, Vol. 48, 2014, pp. 66–83.
35 Ritzinger, U., J. Puchinger, and R. F. Hartl, A Survey on Dynamic and Stochastic Vehicle Routing
36 Problems. *International Journal of Production Research*, Vol. 54, No. 1, 2016, pp. 215–231.
37 Psaraftis, H. N., M. Wen, and C. A. Kontovas, Dynamic Vehicle Routing Problems: Three Decades
38 and Counting. *Networks*, Vol. 67, No. 1, 2016, pp. 3–31.
39 Hickman, J. S., F. Guo, M. C. Camden, N. J. Dunn, and R. J. Hanowski, An Observational Study
40 of the Safety Benefits of Electronic Logging Devices Using Carrier-Collected Data. *Traffic*
41 *Injury Prevention*, Vol. 18, No. 3, 2017, pp. 312–317.
42 Crizzle, A. M., J. Malkin, G. A. Zello, R. Toxopeus, P. Bigelow, and M. Shubair, Impact of
43 Electronic Logging Devices on Fatigue and Work Environment in Canadian Long-Haul
44 Truck Drivers. *Journal of Transport & Health*, Vol. 24, 2022, p. 101295.