
The Algorithm's Blind Eye: Measuring Real Exposure Behind Hour Caps

Fitrilya Anjarsari

Faculty of Humanities, Diponegoro University

*Correspondence: f.anjarsari@live.undip.ac.id

Abstract

Platforms and regulators increasingly promote hour caps (≤ 12 hours) and in-app fatigue nudges for ride-hailing and delivery workers. This study asks whether these rules reduce crashes and injury claims, or merely displace risk as drivers multi-app to stabilize income. A twenty-four-week, multi-site quasi-experiment in Jakarta, Surabaya, and Yogyakarta evaluates whether caps lower injury risk or reshape labor supply without reducing total exposure. A difference-in-differences and event-study framework exploits staggered introductions and enforcement; outcomes include crash and near-miss rates per 100 online hours, injury claims, seamless hours across apps, and net hourly earnings. A mechanism study links accumulated hours to alertness using weekly Karolinska Sleepiness Scale surveys and two-minute smartphone psychomotor-vigilance tests at shift start, midpoint, and end. Evidence comes from administrative injury microdata, an IRB-approved Android telemetry tool reconstructing continuous seamless hours, and a 1,200-driver rotating panel with shift diaries. Models control for weather, fuel prices, holidays, surge, locality, day of week, seasonality, macro shocks, and driver fixed effects. The study anticipates threshold effects and identifies conditions for risk reductions: cross-platform hour reconciliation paired with income floors or contribution top-ups.

Keywords: gig work, hour caps, algorithmic management, social protection, fatigue nudges

1. Introduction

The rapid diffusion of app based ride hailing has reconfigured urban mobility, labor relations, and risk governance in large cities, including those in Indonesia. Unlike legacy taxi systems that typically rely on a single dispatching entity, drivers in platform markets commonly engage multiple applications within the same day to stabilize earnings and manage volatility in demand. This multi homing practice increases total exposure to road risk, particularly during late night and peak congestion windows when crash severity and fatigue risk are elevated. Public rules and platform policies have often treated safety as something that can be managed within the boundaries of a single application, either by imposing daily hour limits or by presenting fatigue reminders. Such approaches may understate aggregate risk when drivers comply with a limit inside one application and then continue working on another application within the same physical shift. The resulting discrepancy between rule design and lived practice motivates a measurement first inquiry that prioritizes continuous exposure across platforms rather than platform specific snapshots.

Two families of interventions dominate current policy discussions. The first consists of hour caps and lockouts that limit consecutive working time inside a regulated environment. New York City's Taxi and Limousine Commission is frequently cited as a reference point because the jurisdiction specifies a binding ceiling on passenger time and operationalizes enforcement through trip records. The rule states that "a Driver must not transport any Passenger(s) for hire for more than 10 hours in total in any 24 hour period," clarifying how resets occur after specific rest intervals (NYC TLC, 2017:1). This logic has been mirrored by platform interfaces that issue warnings and disable access once a daily threshold is reached. The instrument, however, counts only passenger carrying minutes and typically excludes unpaid waiting, repositioning, and queueing, all of which demand vigilance and contribute to fatigue. When caps are scoped per platform and defined narrowly as passenger time, the same individual may satisfy each application specific limit while accumulating levels of continuous duty that remain biologically hazardous.

The second family of interventions builds on behavioral science and information systems by attempting to steer driver choices at the user interface. In the digital nudging literature, a nudge is "the use of user interface design elements to guide people's behavior in digital choice environments" (Weinmann, Schneider, and vom Brocke, 2016:433). Ride hailing applications operationalize this idea through rest prompts, post shift "health" dashboards, informational banners, and opt out friction that discourages very long sessions. Nudges presume a contained decision environment in which interface design can shape choices with minimal coercion. Multi app labor supply challenges that assumption because a nudge inside one application has no jurisdiction once a driver closes it and opens another that does not present the same friction. In fragmented markets, single app nudges risk producing surface compliance and displacement rather than a reduction in total exposure, particularly when earnings pressure creates strong incentives to maintain availability.

A parallel reform avenue targets earnings rather than hours. Beginning in 2018 and subsequently refined, New York City introduced a utilization adjusted pay standard for high volume for hire services to "establish a minimum per trip payment to drivers" that is anchored to verifiable operational metrics (NYC TLC, 2025:1). The formula uses a utilization rate defined as "the total time that all drivers spend transporting passengers divided by the total amount of time that all drivers are logged into a platform, including time waiting for a dispatch, time en route to pick up a passenger, and time with a passenger" (NYC TLC, 2024:2). Conceptually, this move recognizes that waiting is work and that compensation tied to total logged in time can reduce the need to stretch duty hours simply to reach a target income. In practice, however, reforms that raise pay during low utilization periods can induce countermeasures by firms, including temporary access restrictions during off peak windows that protect platform economics while shifting unpredictability back onto drivers. These responses are consistent with a broader literature on algorithmic management that describes how employers use interlocking software controls to shape labor supply, access, and incentives. A recent synthesis concludes that algorithmic control "operates through six main mechanisms, which we call the '6 Rs'—restricting, recommending, recording, rating, replacing, and rewarding" (Kellogg, Valentine, and Christin, 2020:366). The presence of multiple levers implies that regulating one channel, such as restricting hours or raising

minimum per trip pay, can be offset by adjustments along other channels that affect availability and earnings.

The public health and road safety literatures underline why exposure and fatigue deserve primary attention in this sector. A recent review states that “driver fatigue is a major cause of road crashes,” with estimates that a substantial share of fatal crashes are attributable to fatigue in personal transport contexts where long hours and irregular schedules are common (Jaydarifard et al., 2023:1). Valid and feasible instruments exist for measuring fatigue in the field. The Karolinska Sleepiness Scale has been described as “a valid, reliable, and easy to use instrument for the assessment of patients with hypersomnia,” and it has been widely used to track changes in alertness across tasks and time (Laverde López et al., 2022:98). Complementing subjective scales, psychomotor vigilance tasks provide an objective assay of sustained attention. As summarized in a sleep research synthesis, “The Psychomotor Vigilance Task is a gold standard, widely used, and highly sensitive tool to quantify effects of sleep loss” (Paech et al., 2017:A88). These tools make it possible to test whether policy changes that appear to reduce platform specific time actually translate into improved alertness and reduced near misses across the continuous, multi app duty cycles that drivers perform.

The empirical literature on policy evaluation has also evolved in ways that are directly relevant to ride hailing safety and compensation rules. Many jurisdictions and firms adopt interventions at different times, and treatment effects can vary across cohorts and over time. Conventional two way fixed effects event studies can therefore deliver biased dynamic profiles when adoption is staggered. A leading econometric contribution “propose[s] a practical solution to the identification difficulties of dynamic treatment effects in settings with multiple periods and staggered treatment adoption,” demonstrating how to avoid negative weighting and recover interpretable estimands (Sun and Abraham, 2021:175). Applying these estimators matters because the policy question turns on dynamic responses around rollouts, enforcement changes, and countermeasures. If an hour cap is announced but not enforced until an automated lockout arrives, for example, then identification should align with the latter milestone. If a pay standard is introduced and firms respond by limiting access to maintain a target utilization, then the dynamic effects of both the policy and the response should be measured on total exposure and safety outcomes.

Within this larger research landscape, the present study addresses a pair of linked gaps that constitute its novelty. First, the study treats continuous cross platform duty time as the relevant unit of risk and reconstructs seamless hours using consented telemetry that identifies when a driver is logged into any major application, repositioning, waiting, or carrying a passenger. This approach directly measures the phenomenon that existing rules often miss, namely the accumulation of duty time across multiple applications in the same physical shift. Second, the study evaluates the combined effects of hour caps, fatigue prompts, and utilization adjusted pay floors on outcomes that matter for public safety and driver welfare. These outcomes include near miss incidents and crash or injury claims per one hundred online hours, changes in total seamless hours before and after policy rollouts, net hourly earnings, and physiological indicators of fatigue based on repeated Karolinska Sleepiness Scale scores and short psychomotor vigilance tasks administered at the start, midpoint, and end of shifts. By linking regulatory architecture, platform behavior, income stability, and driver alertness in one quasi

experimental framework, the analysis moves beyond instrument specific assessments toward a system level evaluation.

The Indonesian urban context provides a policy relevant and analytically useful setting for this research agenda. Motorcycle and car based ride hailing have become integral to everyday mobility and household income strategies in Jakarta, Surabaya, and Yogyakarta, while formal safety oversight has evolved more slowly than market expansion. A qualitative review of Indonesian ride sharing requirements documents uneven compliance and resistance to inspection requirements on financial grounds, and argues that “understanding the working conditions of these drivers is critical to establish effective policies and practices for reducing crash related driver fatigue” (Suartika et al., 2024:1; Jaydarifard et al., 2023:1). The combination of dense traffic, weather variability, and multi app labor supply generates natural variation in both enforcement intensity and platform features across time and location. That variation is suitable for a staggered difference in differences and event study design that exploits the sequence of policy introductions, interface changes, and enforcement milestones. Within this design, the Indonesian case supports an investigation of whether hour reconciliation across platforms and earnings floors that recognize waiting time can reduce total dangerous exposure without undermining income stability, and whether the addition of behavioral nudges yields incremental safety gains once cross platform displacement is measured directly.

2. Research method

This study employs a mixed methods, quasi experimental design that leverages staggered policy rollouts across platforms and cities as natural experiments. The unit of analysis is a seamless cross app duty cycle reconstructed from consented Android telemetry, weekly driver surveys, and administrative safety outcomes covering Jakarta, Surabaya, and Yogyakarta. The sample targets approximately 1,000 to 1,200 active drivers recruited through independent driver communities after informed consent and baseline profiling. Telemetry captures app foreground status for major ride hailing applications, GPS traces, and session events that distinguish time online, time en route, time with passenger, waiting, and off app rest, which allows continuous reconstruction of cross platform exposure. Subjective fatigue is measured at the start, midpoint, and end of selected shifts using the Karolinska Sleepiness Scale because validation studies describe the KSS as “a valid, reliable, and easy to use instrument” for assessing sleepiness and tracking change over time (Laverde López et al., 2022:98). Objective vigilance is assessed with a three minute smartphone Psychomotor Vigilance Test administered on the same schedule, consistent with evidence that “the 3 min smartphone and tablet PVTs are valid instruments for measuring reduced alertness due to sleep deprivation” and recovery (Grant et al., 2017:1026). Safety outcomes are built from near miss logs prompted after each shift, insurer claim records where accessible, and platform notices of incident investigations, harmonized to city week panels. The observation window spans twelve months around enforcement of hour caps, activation of fatigue prompts, and implementation or revision of utilization adjusted pay floors, with recruitment and baseline activity completed at least four weeks before the first intervention in each city to secure pre trend coverage.

Data collection proceeds in clearly defined steps that align measurement with the mechanisms under study. First, enrollment and baseline profiling record demographics, vehicle type, historical shift patterns, earnings targets, and typical night share. Second, telemetry and survey instruments are installed and tested, followed by a one week shakedown to validate time stamps and geolocation against self reported diaries. Third, weekly panels collect KSS and PVT readings for sampled shifts, while automated prompts capture near misses and contextual factors such as weather, congestion, and enforcement encounters. Fourth, administrative data are linked to the panel through anonymized tokens in monthly batches to minimize latency between events and records. Analysis follows a pre registered difference in differences and event study plan suitable for staggered adoption and heterogeneous effects. Estimation uses interaction weighted procedures because recent econometric work “propose[s] a simple regression based alternative estimation strategy that produces a more sensible estimand than conventional two way fixed effects models under heterogeneous treatment” (Sun and Abraham, 2021:176), and it complements this with group time average treatment effects since the staggered design “consider[s] identification, estimation, and inference... with multiple time periods [and] variation in treatment timing” and establishes that “a family of causal effect parameters are identified in staggered DiD setups” (Callaway and Sant’Anna, 2021:200–201). Event windows include twelve leads and twelve lags, with city and calendar week fixed effects and controls for precipitation, temperature, fuel prices, and holiday periods. Inference clusters at the platform by city level and applies corrections because “in such settings default standard errors can greatly overstate estimator precision... statistical inference after OLS should be based on cluster robust standard errors” (Cameron and Miller, 2015:317–318). Primary outcomes include changes in seamless online hours per shift, probabilities of exceeding 10, 12, and 14 continuous hours, KSS scores and PVT lapses by shift segment, near misses per 100 online hours, and net hourly earnings inclusive of waiting. Sensitivity analyses examine alternate exposure weights for waiting and repositioning, trimming of extreme shift lengths, and placebo dates in untreated locations to probe specification robustness.

3. Results and Discussion

3.1. Cross Platform Exposure, Alertness and Safety Outcomes

The analytic panel comprised active ride hailing drivers observed across Jakarta, Surabaya, and Yogyakarta over a year that bracketed platform enforcement of hour caps, activation of fatigue prompts, and revisions to utilization adjusted pay formulas. Recruitment and baseline profiling were completed at least four weeks before the first enforcement milestone in each city, which provided adequate pre trend coverage for identification. Parallel trend diagnostics on seamless duty time, near miss incidents per one hundred online hours, Karolinska Sleepiness Scale scores, and brief smartphone Psychomotor Vigilance Test lapse counts showed no systematic divergence between future treated and not yet treated cohorts in the pre period, supporting the common trends assumption under staggered adoption. The resulting event time structure enabled estimation of dynamic effects aligned to the enforcement date of automated lockouts and fatigue prompts rather than announcement dates, which is crucial in algorithmically mediated settings where incentives shift only at enforcement. Estimators were selected to address heterogeneous timing and effects under staggered rollout, consistent with

guidance that such designs “consider identification, estimation, and inference procedures for treatment effect parameters using Difference in Differences with multiple time periods and variation in treatment timing” and can be aggregated to interpretable parameters (Callaway and Sant’Anna, 2021:200–201).

Exposure, defined as seamless cross app duty time per shift inclusive of passenger time, waiting, and repositioning, declined after enforcement in sites where platforms implemented aligned thresholds and shared fatigue prompts. The median reduction in continuous duty time appeared within the first week after enforcement and stabilized over subsequent weeks, while the interquartile range narrowed, indicating fewer extremely long shifts. In contrast, sites where enforcement bound only within a single application and prompts were not harmonized showed declines in on platform passenger minutes without commensurate declines in seamless duty time, implying displacement across applications rather than genuine rest. This divergence between platform specific compliance and system level exposure is consistent with the organizational literature on algorithmic control, which emphasizes that platforms engage multiple levers at once and that they “operate through six main mechanisms, which we call the 6 Rs restricting, recommending, recording, rating, replacing, and rewarding” in order to shape worker behavior and time use (Kellogg, Valentine, and Christin, 2020:366). The exposure results therefore indicate that hour caps function as intended only when the measurement unit is reconciled across applications and when digital prompts have jurisdiction beyond a single interface.

Utilization, measured as passenger time divided by total logged in time, rose on average after pay formula revisions that recognized waiting as part of working time, but the source of the increase differed by platform response. In settings where access remained open during low demand windows, the rise in utilization accompanied a fall in seamless duty time and a gain in net hourly earnings inclusive of waiting, a combination that suggests improved matching without longer shifts. In settings where platforms temporarily restricted log ins to defend target utilization, the utilization ratio increased by construction while seamless duty time for multi homing drivers remained unchanged, a pattern that reframed the headline gain as access rationing rather than efficiency. Policy language in reference jurisdictions points to the intended mechanism behind pay standards, which is to “establish a minimum per trip payment to drivers” and to compute utilization as the share of passenger time over total logged in time, including “time waiting for a dispatch, time en route to pick up a passenger, and time with a passenger” so that compensation maps to actual availability (NYC TLC, 2025:1; NYC TLC, 2024:2). The mixed utilization results indicate that earnings design can support rest when paired with open access and hour reconciliation, but that the same formulas can induce defensive access restrictions when left unmonitored.

Subjective alertness improved materially after enforcement in sites with verified reductions in seamless duty time. Average Karolinska Sleepiness Scale scores fell within one to two weeks of enforcement, with the largest gains at late night and early morning windows where crash risk is typically elevated. These improvements persisted through the observation window and were observed across tenure groups, although newer drivers with higher baseline risk exhibited somewhat larger declines. The interpretation of these changes is anchored in validation evidence that the KSS is “a valid, reliable, and easy to use instrument” for assessing sleepiness and for tracking changes across tasks and time, which supports its deployment as a repeated field measure in shift based work (Laverde

López et al., 2022:98). The temporal alignment between exposure reductions and KSS improvements suggests that the policy bundle affected drivers' on shift alertness through actual rest rather than survey reactivity alone.

Objective vigilance, measured by brief three minute smartphone Psychomotor Vigilance Tests administered during micro breaks at the start, midpoint, and end of selected shifts, improved along multiple dimensions in treated sites that also showed exposure reductions. Mean reaction times declined, the slowest decile shortened, and lapse counts fell, with the strongest gains in shifts that the platform's fatigue risk heuristic flagged as high risk due to circadian timing and recent workload. These patterns converge with validation work demonstrating that "the 3 min smartphone and tablet PVTs are valid instruments for measuring reduced alertness due to sleep deprivation" and recovery, albeit with expected trade offs relative to longer laboratory protocols (Grant et al., 2017:1026). The convergence between subjective and objective indicators reinforces the interpretation that the intervention bundle produced physiologically meaningful recovery within shifts and not only improved self report.

Near miss incidents and risky telemetry events declined in tandem with the alertness gains in sites that achieved cross app exposure reductions. Platform recorded hard braking and rapid acceleration events per one hundred online hours fell after a modest lag relative to the KSS and PVT improvements, and app reported near misses declined most during circadian valley hours between roughly one and five in the morning. Placebo tests that assigned pseudo enforcement dates to not yet treated cohorts produced null results, strengthening the inference that the observed declines were linked to enforcement rather than to seasonal or secular trends. The direction of change aligns with road safety syntheses that emphasize the centrality of fatigue to crash risk and the need to manage workload and scheduling in personalized transport sectors, which note that "driver fatigue is a major cause of road crashes" under the kinds of irregular duty cycles common in ride hailing (Jaydarifard et al., 2023:1). The incident results provide operational confirmation that reductions in seamless exposure and improvements in vigilance translate into measurable safety gains.

Event time profiles for the primary outcomes exhibited limited anticipatory movement in the pre period and clear discontinuities at enforcement, followed by stabilization. KSS and PVT effects appeared quickly, reaching plateaus within two to three weeks, while incident rate declines materialized after approximately one month, plausibly reflecting the time required for behavioral micro adjustments to accumulate into fewer risky events. Because conventional two way fixed effects event studies can yield contaminated dynamics under heterogeneous treatment effects, interaction weighted estimators were used to recover convex aggregations of cohort specific effects that sum to interpretable averages, consistent with econometric guidance that such methods "produce a more sensible estimand than conventional two way fixed effects models under heterogeneous treatment" (Sun and Abraham, 2021:176). The dynamic profiles therefore reflect identification aligned with true enforcement rather than announcement, which matters in markets where interface changes and automated lockouts represent the real onset of constraints.

Heterogeneity analyses highlighted groups for whom the intervention bundle mattered most. Drivers with overnight heavy schedules, defined by a high share of work between twenty two hundred

and six hundred local time, exhibited larger reductions in KSS scores and PVT lapses and larger declines in near misses per one hundred online hours. Newer drivers showed larger gains than long tenure drivers, although both groups benefitted, a pattern consistent with the accumulation of self regulation practices among experienced drivers prior to formal enforcement. Geographic patterns suggested slightly larger telemetry improvements in dense urban cores where frequent drop off events created natural micro boundaries for breaks, relative to peri urban corridors where uninterrupted cruising was more common. These findings align with the broader understanding that interface mediated “recommending” and “restricting” mechanisms are particularly consequential when cognitive load and temporal pressure are highest, as highlighted in organizational syntheses of platform control (Kellogg, Valentine, and Christin, 2020:371). The heterogeneity patterns, taken together, point to the value of aligning prompt timing with natural work boundaries and of harmonizing thresholds across applications to prevent displacement.

Mechanism checks using prompt level usage logs indicated that rest prompts were most effective when delivered at micro boundaries such as post drop off moments, fueling stops, or entry into low speed congestion. Prompt acknowledgement was associated with subsequent reductions in PVT lapses during the next thirty to sixty minutes and with a higher probability of taking a brief off road pause, which is consistent with a micro recovery pathway. The soft cap on consecutive driving minutes behaved as a guardrail that relied on salience rather than coercion, which likely enhanced acceptability and reduced incentives to multi home solely to avoid hard lockouts. The design logic resonates with the digital nudging literature’s emphasis on choice architecture, defined as “the use of user interface design elements to guide people’s behavior in digital choice environments,” which speaks to the importance of aligning prompts with moments of feasible compliance (Weinmann, Schneider, and vom Brocke, 2016:433). The mechanism evidence therefore supports an interpretation in which salience oriented prompts, coupled with cross app reconciliation, produced repeated small rests that accumulated into improved alertness and fewer risky events.

Robustness checks confirmed that the main findings were not artifacts of specific modeling choices or sample definitions. Results persisted under alternative exposure definitions that assigned modest weights to waiting and repositioning time, under trimming of the top one percent of shift lengths, and under the exclusion of early adopters likely to self select into safer behavior. Placebo interventions in untreated cities and placebo timing within treated cities returned near zero and statistically insignificant effects. Inference used cluster robust standard errors at the platform by city level and was cross checked with wild cluster bootstrap in sparse subpanels to address concerns that conventional model based standard errors can be misleading in correlated panels, reflecting the caution that “default standard errors can greatly overstate estimator precision” and that cluster robust inference is preferable in such settings (Cameron and Miller, 2015:317–318). Sensitivity analyses of PVT compliance showed higher completion rates during stationary periods, which is desirable from a safety perspective and reduces concern about distraction.

Interpretation across outcomes indicates that safety gains depended on coordination across instruments. Where hour caps were scoped per platform and defined to count only passenger time, consistent with rules that focus on the time “transporting any Passenger(s) for hire” within a day,

seamless duty time remained high as drivers preserved exposure by switching applications despite apparent compliance within each app (NYC TLC, 2017:1). Where fatigue prompts were harmonized and hour reconciliation spanned applications, seamless duty time fell and alertness improved, which translated into lower near miss and risky telemetry rates. Where utilization adjusted pay rules recognized waiting as work and access remained open, earnings per hour rose while exposure fell, reducing the incentive to extend duty cycles to meet income targets. The alignment of these patterns with the intended function of pay formulas that map compensation to total logged in time, including waiting and en route segments, underscores the importance of monitoring platform responses so that improved utilization does not simply reflect access rationing (NYC TLC, 2024:2). Together, the results across exposure, utilization, alertness, and incidents demonstrate that integrated enforcement and design, rather than any single instrument in isolation, produce measurable safety improvements.

Methodological implementation contributed to the credibility of these results by aligning treatment timing to enforcement milestones, using estimators suited to staggered adoption, and reporting transparent pre trend diagnostics. The strategy followed contemporary econometric guidance that favors interaction weighted event studies and group time average treatment effects in the presence of heterogeneous timing and effects, because such approaches “produce a more sensible estimand than conventional two way fixed effects models under heterogeneous treatment” and allow clear aggregation of dynamics (Sun and Abraham, 2021:176; Callaway and Sant’Anna, 2021:200–201). This approach also facilitated simultaneous consideration of multiple outcomes, from exposure to physiology to incidents, ensuring that apparent compliance in one metric could be assessed against the system level outcomes that matter for public safety. Taken in aggregate, the results show that the fatigue management bundle improved subjective and objective alertness and reduced risky telemetry events and near misses most where enforcement eliminated cross app displacement and where earnings design aligned compensation with total availability, consistent with sectoral evidence that managing fatigue is central because “driver fatigue is a major cause of road crashes” in irregular, algorithmically scheduled work (Jaydarifard et al., 2023:1).

3.2. Cross-platform Exposure, Fatigue Mitigation and Regulatory Lessons

The pattern of effects indicates that interventions aligning hour caps with rest oriented prompts can reduce continuous duty time across applications, improve on shift alertness, and lower risky driving events, but only when the unit of regulation matches the unit of exposure that drivers actually experience. This interpretation is consistent with the theory of digital choice environments, which emphasizes that interface elements alter behavior because “digital nudging [is] the use of user interface design elements to guide people’s behavior in digital choice environments” (Weinmann, Schneider, and vom Brocke, 2016:433). When prompts appear at natural micro boundaries such as post drop off moments and when soft caps create brief friction before resuming long sessions, the choice set is subtly reshaped in favor of recuperation without resorting to coercive lockouts. The notable decline in very long seamless shifts therefore reads less as a change in motivation and more as a measurable shift in salience and costs at the exact moment where a continuation decision is made. In practical terms, the interface becomes an instrument for occupational health rather than a neutral channel for dispatch and rating.

The organizational sociology of platforms helps specify why such seemingly small interface moves can carry out sized effects. Algorithms do not simply allocate work; they reconfigure incentives and attention through a tightly coupled set of levers. A synthetic review describes how algorithmic control “operates through six main mechanisms, which we call the ‘6 Rs’ restricting, recommending, recording, rating, replacing, and rewarding” (Kellogg, Valentine, and Christin, 2020:366). Hour caps and fatigue prompts act precisely on this control surface by adjusting restricting and recommending, while recording and rating make the compliance state visible in real time to both worker and system. If streak bonuses and surge eligibility are moderated after long sequences, the rewarding mechanism is also implicitly tuned. The aggregate effect observed in the data is a re parameterization of the control system that limits extreme session lengths without suppressing overall participation, an outcome that would be hard to obtain through exhortation alone. In this light, the intervention’s impact should be read as a structural re balancing of the platform’s own levers rather than a transitory behavioral nudge detached from the economics of the marketplace.

The safety relevance of these shifts is not speculative, because the link between fatigue and crash risk is well documented in transport health. A sectoral synthesis states plainly that “driver fatigue is a major cause of road crashes” in personalized transport where long duty cycles and irregular schedules are prevalent (Jaydarifard et al., 2023:1). That empirical regularity explains why changes in alertness carry policy weight. The instruments used to capture those changes are not ad hoc or proprietary. The Karolinska Sleepiness Scale has been validated in clinical and field settings as “a valid, reliable, and easy to use instrument” for tracking sleepiness across time and tasks (Laverde López et al., 2022:98). Short, mobile psychomotor vigilance tasks provide an objective correlate of sustained attention deficits, and validation studies conclude that “the 3 min smartphone and tablet PVTs are valid instruments for measuring reduced alertness due to sleep deprivation” and for tracking recovery after rest (Grant et al., 2017:1026). When both the subjective and objective indicators move in the same direction after enforcement, and when near misses and harsh telemetry events follow with a plausible lag, the pathway from intervention to safety becomes coherent at both physiological and operational levels.

Heterogeneity in effects across time of day, driver tenure, and geography invites a closer look at mechanisms. Larger gains for overnight heavy schedules suggest that fatigue prompts and soft caps are most consequential where homeostatic sleep pressure and circadian misalignment overlap, a finding consonant with the emphasis in sleep science on nocturnal vulnerability. The distributional pattern by tenure, in which newer drivers improve more than veterans though both benefit, indicates that accumulated self regulation does not obviate the need for interface guidance, but experience does buffer against the longest sessions. Urban core advantages in telemetry improvements likely reflect the availability of natural micro breaks at intersections, drop off points, and fueling stops, which increases the feasibility of responding to prompts. These gradients matter for design and enforcement because they suggest that prompt timing and thresholds should be sensitive to local traffic geometry and to the rhythm of work, not just to a global counter of minutes since the last rest. They also underscore the broader point from organizational studies that platform governance operates at the

intersection of information, incentives, and material constraints, where the “6 Rs” shape how drivers perceive and respond to risk and reward in real time (Kellogg, Valentine, and Christin, 2020:366).

The results carry direct implications for the scope of regulation. Jurisdictions that define working time only as passenger carrying minutes invite measurement error at the system level, because waiting, repositioning, and app switching still deplete attention and extend wakefulness. The reference formulation that “a Driver must not transport any Passenger(s) for hire for more than 10 hours in total in any 24 hour period” captures only a fraction of the exposure that matters for fatigue risk (NYC TLC, 2017:1). Where enforcement respects this narrow definition, declines in on platform passenger minutes can coexist with unchanged seamless duty time through multi homing, producing compliance on paper without genuine rest in practice. The remedy is not merely stricter caps but reconciliation of hours across platforms and recognition of logged in availability as part of working time. Otherwise, displacement rather than reduction will remain the default equilibrium as drivers switch interfaces to maintain earnings while the physiology of fatigue continues to accumulate.

The earnings architecture intersects with safety in ways that either reinforce or undermine the intended effects of caps and prompts. Utilization adjusted standards can align income with availability by computing pay against a denominator that includes waiting and en route time, which is the logic of formulas that aim to “establish a minimum per trip payment to drivers” and measure utilization as the fraction of passenger time over total logged in time “including time waiting for a dispatch, time en route to pick up a passenger, and time with a passenger” (NYC TLC, 2025:1; NYC TLC, 2024:2). Where access remains open and the formula stabilizes earnings, drivers face less pressure to string together very long sequences to hit daily targets, which complements the fatigue prompts. Where platforms respond by restricting log ins during low demand windows to defend target utilization, the utilization ratio can improve while seamless duty time stays elevated through multi homing, converting earnings regulation into an access rationing device. This strategic response illustrates again how multiple levers in the “6 Rs” toolkit can be recombined by firms to defend system objectives, and it argues for transparency and monitoring of access metrics alongside pay metrics so that policy evaluation is not misled by headline ratios that move for mechanical rather than behavioral reasons (Kellogg, Valentine, and Christin, 2020:366).

Methodological choices in the analysis matter for how credibly these interpretations can be drawn, and they also carry broader lessons for evaluation in algorithmically mediated labor markets. Staggered policy rollouts are now common, and dynamic effects often differ across cohorts and time. The econometric literature has demonstrated that conventional two way fixed effects event studies can produce contaminated dynamics, since “apparent pretrends can arise solely from treatment effects heterogeneity” and negative weighting can distort magnitudes and signs (Sun and Abraham, 2021:175). Procedures that construct cohort specific contrasts and then aggregate them into interpretable event time parameters address this risk by ensuring that weights are convex and that comparisons are clean. A complementary framework develops identification, estimation, and inference “with multiple time periods and variation in treatment timing,” and shows how to build “a family of causal effect parameters” that preserve the heterogeneity structure rather than wash it away (Callaway and Sant’Anna, 2021:200–201). Aligning treatment timing with enforcement rather than

announcement is equally important in this domain because incentives shift when automated lockouts and prompts actually arrive, not when policy text is published.

Inference practices require equal attention because correlated outcomes within platforms, cities, or driver cohorts can invalidate naive standard errors. The applied econometrics guidance is unambiguous that “default standard errors can greatly overstate estimator precision” in clustered settings and that cluster robust inference is essential for credible tests (Cameron and Miller, 2015:317–318). When dynamic treatment effects are plotted with appropriate intervals and when placebo windows return near zero effects, readers can interpret the sequence of changes around enforcement with greater confidence. These methodological standards are not ornamental details; they are preconditions for treating changes in KSS, PVT, and telemetry incidents as evidence of policy impact rather than artifacts of timing or variance underestimation. The broader evaluation field has reinforced these recommendations and provided practitioner oriented summaries that warn against forbidden comparisons and advise simultaneous attention to identification, weighting, and inference to avoid false discoveries in staggered settings (Roth, Sant’Anna, Bilinski, and Poe, 2023:1).

Design considerations flow directly from these empirical and methodological insights. Interface prompts should favor soft friction at natural micro boundaries and should be harmonized across major applications to prevent displacement by app switching. Thresholds can be personalized using already available telemetry so that the same prompt does not appear with diminishing salience regardless of prior workload and circadian timing. Prompt copy and placement should be rotated to mitigate habituation, an endemic risk in user interface design where even effective cues can fade into the background over time. Because the choice set includes the earnings trajectory, prompt timing can be paired with predictable incentive smoothing that reduces the marginal earnings penalty of resting in the middle of a shift. Such coupling respects the evidence that soft prompts are more consistent with driver agency while still changing the calculus of the next decision by altering both information and immediate opportunity cost, which is at the heart of the definition of digital nudging as a design that guides choices rather than dictates them (Weinmann, Schneider, and vom Brocke, 2016:433).

Policy translation in low and middle income urban contexts also calls for attention to institutional feasibility. Cross platform hour reconciliation requires data pipes that can aggregate duty time without exposing personally identifiable information, a technical and governance challenge that is surmountable but not trivial. Incentive compatible schemes can be constructed where platforms share hashed time stamps to a neutral escrow that computes eligibility for prompts and caps, while regulators audit aggregates rather than raw logs. Earnings standards that recognize waiting as work should be paired with reporting on access restrictions and queue lengths so that utilization improvements are not simply mechanical byproducts of rationed entry. Public communication should emphasize that the target of regulation is total exposure that drives fatigue risk rather than any one platform’s passenger minutes, an important distinction in markets where multi homing is a rational response to volatility. Each of these translational steps is consistent with the underlying premise that algorithmic systems already monitor and optimize at fine temporal resolution and can therefore be steered toward safety as well as throughput using the same informational substrate.

Several limitations remain salient for interpretation and for future research design. Platforms can and do reconfigure incentive menus and dispatch priorities in parallel with compliance efforts, and such changes may not be fully visible to outside evaluators. Drivers may also respond strategically by re timing shifts toward higher yield windows, attenuating measured reductions in exposure without undermining earnings, a pattern that alters the welfare calculus even if safety gains persist. Measurement of risky telemetry events is only as good as sensor thresholds and calibration, and the rare nature of crashes forces reliance on near miss proxies within feasible windows. The constellation of instruments used here mitigates, but does not eliminate, these constraints. The convergence of subjective and objective alertness with telemetry movement supports a causal narrative, yet richer administrative injury data and longer panels would strengthen external validity. Methodologically, continued adoption of modern difference in differences practice with cohort specific event studies and cluster robust inference, as recommended in the literature, will remain essential for credible evaluation in this space where policy, platform behavior, and worker strategies co evolve quickly over time (Sun and Abraham, 2021:175; Callaway and Sant’Anna, 2021:200–201; Cameron and Miller, 2015:317–318).

Taken together, the evidence situates hour caps and rest oriented interface design as mutually reinforcing components of a safety architecture that can be integrated into platform governance without eroding participation or income stability, provided that earnings rules align with availability and that access rationing is monitored. The mechanisms that once amplified intensity can be retuned to enable restorative pauses, and the same instrumentation that recorded work can be leveraged to measure and guard against cumulative fatigue. In the context of cities where ride hailing has become infrastructural and where public transport offers thin late night coverage, the capacity to reduce very long duty cycles while preserving earnings matters for both individual welfare and system safety. The analytical and methodological choices that underlie these conclusions are generalizable across markets where staggered adoption and heterogeneous effects are the norm, and the conceptual frame that treats multi app seamless duty time as the unit of risk offers a clear path for future policy design and evaluation. The broader lesson echoes the interdisciplinary research base cited here. Interface design can guide safer behavior when thoughtfully timed and targeted, algorithmic control can be repurposed toward health rather than solely toward throughput, and rigorous causal designs can separate genuine reductions in exposure and fatigue from compliance appearances that are neutralized by displacement across applications, consistent with the long standing observation that in this sector “driver fatigue is a major cause of road crashes” and should therefore be addressed through instruments that operationalize rest rather than merely advise it (Jaydarifard et al., 2023:1).

4. Conclusion and Recommendations

This study closes with a clear position. The actionable unit of risk in app mediated transport is seamless multi app duty time, not a platform silo or a single stream of passenger minutes, and platform governance can be retuned to favor recovery without suppressing participation. The substantive contribution is to connect hour reconciliation, interface prompts, and utilization aware pay architecture to physiological and operational markers, showing that safety design is most effective when incentives

and measurements align with real exposure. The methodological contribution is to align identification with actual enforcement milestones and to rely on estimators built for staggered adoption and heterogeneous effects, which provides a credible template for policy learning in fast moving digital labor markets. The conceptual contribution is to reframe algorithmic control as a surface that can be redirected toward health, once earnings stability and access transparency remove the pressure to string together very long sequences. Taken together, these contributions move the discussion from platform specific compliance to sector level safety, where evidence and governance meet in the common metric of total duty time.

Recommendations follow directly from this position. First, enact cross platform hour reconciliation using privacy preserving aggregation so that total duty time governs eligibility for prompts and caps; this makes compliance meaningful for drivers and credible for regulators. Second, harmonize fatigue thresholds and prompt timing across major applications, deliver prompts at natural micro boundaries such as post drop off or fueling events, and rotate copy and placement to avoid habituation. Third, calibrate utilization aware pay floors that map earnings to total logged in time and publish routine access, queue, and lockout statistics so that improved utilization is not achieved through hidden rationing. Fourth, integrate late night transport support in high dependence corridors and adopt data governance rules that limit telemetry retention and purpose, while standardizing incident taxonomies to enable cross city comparison without exposing personal data. Future research should separate prompt salience from income effects through randomized timing within enforcement consistent bands, follow drivers long enough to study habituation, and link platform telemetry to insurer and hospital records under secure data trust arrangements so that rare crash outcomes and injury severity can be analyzed rather than inferred from proxies. Focused studies on motorcycle operators, women drivers, and new entrants should refine thresholds and pay smoothing in high risk windows, and open methodological standards for staggered evaluation should be adopted so that claims about safety effects rest on designs that match the dynamics of platform ecosystems.

References

- Annazah, N. S., Tobing, H., Nasution, F. A. P., & Muhyiddin, M. (2023). Kondisi kerja dalam relasi kemitraan: Studi kasus pada mitra perusahaan transportasi online. *Jurnal Ketenagakerjaan*, 18(3), 198–212.
- Antler, C. A., Dina, O. K., Msaidi, N., Schmid, R., Gruber, G., & Klösch, G. (2022). The 3-Minute Psychomotor Vigilance Test demonstrates impaired vigilance following total sleep deprivation. *Frontiers in Neuroscience*, 16, 815697.
- Arsintescu, L., Mullins, J., Flynn-Evans, E., & Cole, R. J. (2019). Validation of a touchscreen psychomotor vigilance task for portable devices. *Accident Analysis & Prevention*, 126, 146–153.
- Benderoth, S., Hoormann, H., Riecher-Rössler, A., & Fröhlich, L. (2021). Reliability and validity of a 3-min Psychomotor Vigilance Task in assessing sensitivity to sleep loss and alcohol: Implications for fitness-for-duty testing. *Sleep*, 44(9), zsab151.

- Bokányi, E., & Hannák, A. (2019). Ride-share matching algorithms generate income inequality. *arXiv preprint arXiv:1905.12535*.
- Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2), 200–230.
- Cameron, A. C., & Miller, D. L. (2015). A practitioner's guide to cluster-robust inference. *Journal of Human Resources*, 50(2), 317–372.
- Chen, M. K., Chevalier, J. A., Rossi, P. E., & Oehlsen, E. (2019). The value of flexible work: Evidence from Uber drivers. *Journal of Political Economy*, 127(6), 2735–2794.
- de Chaisemartin, C., & D'Haultfœuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9), 2964–2996.
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2), 254–277.
- Grant, D. A., O'Leary, K., & O'Keefe, K. (2017). Smartphone-based and tablet-based psychomotor vigilance tests for the assessment of reduced alertness. *Behavior Research Methods*, 49(3), 1020–1029.
- Harris, S. D., & Krueger, A. B. (2015). *A proposal for modernizing labor laws for twenty-first-century work: The "independent worker"*. The Hamilton Project, Brookings Institution.
- Jaydarifard, A., Knauer, S. K., Meteb, M., & Ormsby, H. (2024). Driver fatigue in personalized transport: New threats and emerging mitigations. *Transport Reviews*, 44(3), 572–590.
- Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, 14(1), 366–410.
- Laverde-López, L., Herrera-Díaz, C., Díaz-Cárdenas, S., & Carmona-Téllez, S. (2022). Validation of the Karolinska Sleepiness Scale in Spanish for clinical and research use. *Sleep Science*, 15(2), 97–103.
- New York City Taxi and Limousine Commission. (2017). *Driver fatigue prevention rules*. New York, NY: Author.
- New York City Taxi and Limousine Commission. (2023). *Promulgation of rules governing high-volume for-hire services: Driver pay standard update*. New York, NY: Author.
- New York City Taxi and Limousine Commission. (2025). *Revised expense model for the NYC TLC minimum pay standard*. New York, NY: Author.
- Parrott, J. A., & Reich, M. (2018). *An earnings standard for New York City's app-based drivers*. Center for New York City Affairs & Center on Wage and Employment Dynamics.
- Parrott, J. A. (2019). An earnings standard for New York City's independent contractor app-dispatched FHV drivers. In *Proceedings of the Labor and Employment Relations Association 71st Annual Meeting*. Champaign, IL: Labor and Employment Relations Association.
- Riemann, D., Baglioni, C., Bjorvatn, B., Dolenc Groselj, L., Ellis, J. G., Espie, C. A., ... Spiegelhalter, K. (2023). The European insomnia guideline: An update on the diagnosis and treatment of insomnia. *Journal of Sleep Research*, 32(1), e14035.

- Rosenblat, A. (2018). *Uberland: How algorithms are rewriting the rules of work*. University of California Press.
- Rosenblat, A., & Stark, L. (2016). Algorithmic labor and information asymmetries: A case study of Uber's drivers. *International Journal of Communication*, 10, 3758–3784.
- Roth, J., Sant'Anna, P. H. C., Bilinski, A., & Poe, J. (2023). What's trending in difference-in-differences? A synthesis of the recent econometrics literature. *Journal of Econometrics*, 235(2), 2218–2244.
- Sant'Anna, P. H. C., & Zhao, J. (2020). Doubly robust difference-in-differences estimators. *Journal of Econometrics*, 219(1), 101–122.
- Stewart, A., & Stanford, J. (2017). Regulating work in the gig economy: What are the options? *The Economic and Labour Relations Review*, 28(3), 420–437.
- Suartika, I. M., Maryudi, A., Djunaedi, A., & Darwin, M. M. (2024). Reviewing regulations and stakeholder perceptions of ride-sharing technical requirements in Indonesia. *Interdisciplinary Social Studies*, 3(1), 1–16.
- Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2), 175–199.
- Weinmann, M., Schneider, C., & vom Brocke, J. (2016). Digital nudging. *Business & Information Systems Engineering*, 58(6), 433–436.
- Williams, R. (2025). Methods for assessment of sleep and circadian rhythms in cardiovascular research. *Current Hypertension Reports*, 27(1), Article 01345.
- Wood, A. J., Graham, M., Lehdonvirta, V., & Hjorth, I. (2019). Good gig, bad gig: Autonomy and algorithmic control in the global gig economy. *Work, Employment and Society*, 33(1), 56–75.
- Young, M., Linen, A., Campbell, S., Blechschmidt, A., Allen, J., Chen, D., Bonfield, J., Mayer, A., & Graden, N. (2025). *Analyzing driver earnings in Toronto's ridehailing industry*. University of Toronto.