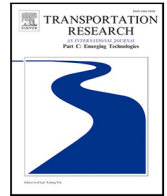Contents lists available at ScienceDirect

Transportation Research Part C

journal homepage: www.elsevier.com/locate/trc

Managing lane-changing of algorithm-assisted drivers

Mihalis G. Markakis [a,*], Kalyan Talluri [b], Dmitrii Tikhonenko [b]

[a] *IESE Business School, University of Navarra, Spain*
[b] *Imperial College Business School, United Kingdom of Great Britain and Northern Ireland*

ARTICLE INFO

ABSTRACT

Theoretical models of vehicular traffic ascribe the fundamental cause of velocity oscillations and stop-and-go waves to suboptimal or unpredictable human driving behavior. In this paper we ask: if vehicles were controlled or assisted by algorithms, and hence driven "optimally," would these phenomena simply go away? If they do not, how should a regulator manage algorithm-assisted vehicular traffic for a smooth flow? We study these questions in the context of a mandatory lane-changing scenario from the perspective of an algorithm-assisted driver on a curtailed lane that has to merge to an adjacent free lane with a relatively dense platoon. In a stylized model of algorithm-assisted driving, we liken the blocked-lane driver to a rational self-interested agent, whose objective is to minimize her expected travel time through the blockage, deciding (a) at what velocity to move, and (b) whether to merge to the free lane if an adequate gap arises. Moving at higher velocities reduces travel time, but also reduces the probability of finding a large enough gap to merge. We analyze the problem via dynamic programming, and we show that the optimal policy has a surprising structure: in the presence of uncertainty on adequate-sized gaps in the target lane, it may be optimal for the blocked-lane driver, in certain parameter regimes, to oscillate between high and low velocities while attempting to merge. Hence, traffic oscillations can arise not just due to suboptimal or unpredictable human driving behavior, but also endogenously, as the outcome of a driver's rational, optimizing behavior. We provide theoretical support for this finding by drawing similarities to bang–bang control. As velocity oscillations are known to be detrimental to a smooth traffic flow, we provide sufficient conditions such that traffic oscillations, due to such optimizing behavior, do not arise. Finally, we investigate the fundamental flow-density and travel time-density diagrams through traffic micro-simulations performed in SUMO. We establish that the proposed approach exhibits consistently near-optimal performance, in a broad variety of traffic conditions.

1. Introduction

Smooth vehicular traffic flow can reduce environmental pollution, improve quality of life for commuters, and save billions of dollars in lost productivity. The nature of road traffic, however, poses significant modeling and technical challenges: it is highly dynamic and hard to predict, with a large number of autonomous self-interested agents of varying skills and attitudes, with the decisions of one agent potentially affecting the actions of many others. Moreover, vehicular traffic is generally controlled only by coarse static rules and guidelines that, by necessity, have to be applicable to a broad spectrum of driving conditions.

Despite many decades of sustained research by various communities, a number of questions and challenges in vehicular traffic control remain, as evident from the prevalence of traffic jams and delays to this day (Treiber and Kesting (2018)).

* Corresponding author.
E-mail addresses: MMarkakis@iese.edu (M.G. Markakis), kalyan.talluri@imperial.ac.uk (K. Talluri), d.tikhonenko@imperial.ac.uk (D. Tikhonenko).

This paper focuses on (mandatory) lane changes for *algorithmic* drivers and explores the phenomenon of traffic oscillations, which plays an important role in creating traffic jams. For our purposes, algorithm-assisted driving signifies: (a) the algorithm is assumed to have a well-defined and programmed objective function that it optimizes, in our case minimizing travel time; and (b) the optimization is based on a model of information and behavior of other drivers, in our case a simple stochastic model that can be calibrated based on empirical data and past traffic studies.

Many macroscopic as well as microscopic models of vehicular traffic use exogenous stochastic elements to emulate typical but inexplicable human driving behavior, such as sudden or excessive braking or acceleration, or reacting with lags; for instance, see Section G of the survey paper of Helbing (2001) on a model with a random variable to capture human driving errors. In simulations, these stochastic errors may then propagate upstream and make the traffic flow unstable, occasionally even causing it to come to a complete halt for seemingly no reason. By implication, the causal factor for stop-and-go waves and traffic oscillations is the random nature of human driving. This naturally raises the question that we address in this paper: *Would traffic oscillations simply go away when vehicles are driven by algorithms?*

This question has taken on renewed urgency with the advent of sensors, 5G-based vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communication, and driverless car technologies, which enable significant information exchange as well as more precise vehicle control, and allow algorithmic decision-making that is based on clear objectives. At the time of this writing, major car manufacturers such as Audi, Honda, Mercedes-Benz, Tesla and BMW have achieved vehicles with Level 3 Autonomy (conditional driving automation), and promise to deliver Level 4 Autonomy (high driving automation) within the next few years. Traffic regulators are expected to follow suit, albeit with an understandable time lag; for instance, Germany is expected to lead the adoption of Level-3 infrastructure in Europe in the near future, followed by France and the UK. These technological developments raise the need for better understanding, via models and analysis, of optimal policies that take drivers' objectives and self-interested nature into consideration.

We study these issues in the context of a specific traffic situation: mandatory lane-changing with "mixed traffic", i.e., an algorithm-assisted driver on a curtailed lane that has to merge to an adjacent free lane with a relatively dense platoon of human drivers. In a stylized model of algorithm-assisted driving, we liken the blocked-lane driver to a rational self-interested agent whose objective is to minimize her expected travel time through the blockage, deciding (a) at what velocity to move, and (b) whether to merge to the free lane if an adequate gap presents itself. Notably, the blocked-lane driver has only a *stochastic* view of the state of the free lane. Indeed, even if she knows the exact locations of merging opportunities at the time of her planning decisions, the gaps may not be there, or may be of insufficient size, when she actually gets there due to the driving decisions of the free-lane drivers.

What makes the problem interesting is that moving at higher velocities on the blocked lane reduces travel time in the short run but requires larger gaps to merge to the denser free lane as suggested by gap-acceptance theories of merging. Consequently, blocked-lane drivers are trading off faster travel time at the current stage with an increased risk of not being able to merge before reaching the end of the road segment, where they may be forced to come to a complete halt.

The contributions of this paper are as follows:

1. On the modeling side, we provide a novel Dynamic Programming formulation for an algorithm-assisted vehicle that has to merge. Our model can be viewed as an extension to the classical "parking problem" in Operations Research, and although a stylized representation of reality, captures the salient features of the situation under study.
2. We characterize the optimal policy and highlight its structure: in the presence of merging uncertainty, it may be optimal for a driver, in certain parameter regimes, to oscillate between high and low velocities while attempting to merge. So, algorithm-assisted driving based on a model and (global) optimization of an objective function may lead to oscillatory traffic patterns. Hence, our main insight is that traffic oscillations can arise endogenously, even with rational driving behavior. Algorithms may eliminate some sources of velocity oscillations, associated with human driving behavior, but may introduce some new ones too.
3. We shed further light into the mathematical structure of the problem by arguing, in the context of a more abstract version of our benchmark model with continuous state and action spaces, that the optimal policy is an instance of bang–bang control. This analysis also reinforces our main finding, that regions of maximum acceleration (thus, high velocity) and deceleration (thus, low velocity) are a structural characteristic of the problem, and not the artifact of modeling choices.
4. On the normative side, we show how a traffic regulator can set velocity limits, so that velocity oscillations do not arise at optimality. This would lead to a safer, smoother, and more efficient traffic flow.
5. We construct a realistic representation of the mandatory merging scenario under study using the SUMO traffic micro-simulator; see Lopez et al. (2018). Through simulations, we investigate the fundamental flow-density and travel time-density diagrams of the proposed DP-based policy against various benchmarks. Our findings establish that the policy provides a principled way to achieve good performance in a broad variety of traffic conditions, whereas the different benchmarks such as early/late merge perform well only under certain traffic conditions.

Our research illustrates some new dimensions that algorithm-assisted drivers can bring to vehicular traffic. Human drivers, in local traffic maneuvers, can generally perform only *local* optimization in a heuristic way, based on their immediate observed surrounding vehicles and conditions. Algorithm-assisted drivers, on the other hand, are capable of making *globally* optimal decisions based on models of information and other drivers' behavior, due to their ability to both exchange as well as process significant amounts of information in real-time. Such global optimization, at an individual-driver level, can in fact lead to undesirable

phenomena that were expected to be eliminated by algorithmic driving. Hence, there is a need for altruistic or cooperative behavior – such as courtesy yielding – that has to be either incentivized or imposed by a traffic regulator.[1]

We note that our approach is fundamentally different from both empirical and theoretical studies in the traffic literature. One, theoretical traffic models, say differential equations that link macroscopic measures such as density and flow, may fit observed data well but do not pinpoint precise causal reasons for a particular phenomenon as does our model. Second, the very simplicity of our model differentiates it from complex microscopic simulations that are not amenable to a provable conclusion based on analysis; such simulation set-ups do capture many features of human driving behavior, but often have too many confounding factors to pinpoint fundamental causes. Finally, as already mentioned, our focus is on algorithm-assisted driving rather than human. One implication is that data on human driving, and the myriad empirical studies based on them, while certainly relevant, cannot be used to infer phenomena potentially created due to global optimization by algorithm-assisted drivers.

The remainder of the paper is organized as follows. In Section 2, we survey the related literature, focusing primarily on articles that research lane-changing. In Section 3, we present a stylized discrete-stage and velocity model of lane-changing in forced-merge scenarios, and show that velocity oscillations can occur at optimality. In Section 4, we show how a central planner can set velocity limits to eliminate these kind of oscillations. In Section 5, we argue that the optimal policy is an instance of bang–bang control. We conclude the paper in Section 6 by investigating the fundamental flow-density and travel time-density diagrams of the proposed DP policy in more realistic environments, through SUMO traffic micro-simulations. The appendices provide the proofs of our analytical results; an analysis of the NGSIM data set, which shows that oscillations like the ones we study here already exist in human driving; details on our simulation set-up; and extensions to our benchmark model.

2. Literature on lane-changing

Traffic modeling follows either a macroscopic approach, studying traffic at the population/fluid level, or a microscopic approach, studying the behavior of individual drivers. Both strands of literature are too vast, and the models are too many to even attempt to summarize here. Instead, we refer the reader to the survey papers of Helbing (2001) and Chowdhury et al. (2000), from a macroscopic and a microscopic point of view respectively. In the rest of this section, we emphasize the research relevant to lane-changing.

Over the last two decades, lane-changing has gained prominence in the traffic literature. It is increasingly recognized as one of the principal reasons for creating disturbances in traffic flows, as well as oscillations leading to stop-and-go waves in traffic patterns (Zheng et al. (2010)).

A notable strand of literature lies in the intersection of Traffic Engineering and Computer Science, where the details and the effects of lane-changing are studied through *micro-simulations*. Chang and Kao (1991), Nagel and Schreckenberg (1992), Fritzsche (1994), Rickert et al. (1996), Wagner et al. (1997) and Barlovic et al. (1998, 1999) use Cellular-Automata models, setting driver-level rules for lane-changing on multi-lane highways. More sophisticated micro-simulation environments are proposed in Yang and Koutsopoulos (1996) and Flotterod et al. (2011). From this literature, let us further highlight Ebersbach and Schneider (2004) and Han and Ko (2012), which focus on merging into a highway from a blocked lane and an on-ramp respectively. Similar in spirit is the approach in Tobita et al. (2012), which studies merging and bifurcations in a two-lane highway. We note that the above literature models and simulates the effects of lane-changing of human drivers with local information and behavior. In contrast, our focus is on algorithm-assisted drivers, whose decisions are based on the global optimization of a clearly defined objective.

The Traffic Engineering community has proposed many lane-changing *models*, estimating their parameters from real data or sophisticated micro-simulation models, e.g., Gipps (1986), Kita (1999), Daganzo (2002a), Daganzo (2002b), Hidas (2002), Jin (2010), Laval and Daganzo (2006), Kondyli and Elefteriadou (2011), Zhang et al. (2012), Keyvan-Ekbatani et al. (2016) and Ngoduy et al. (2019). Choudhury and Ben-Akiva (2013) and Kesting et al. (2007) propose sophisticated merging decision structures and regimes, including potential cooperation with cars in the target lane. A number of papers model velocity and accelerations as a part of merging maneuver. A typical approach is to use functional forms that depend on certain external conditions, such as lead gap and velocities of other actors, and then estimate the coefficients on the real data. Notable works of this nature are Choudhury et al. (2009) and Wan et al. (2017). However, the velocity decisions are not endogenous, in that are a result of an optimal decision process, which differentiates these studies from our work. Finally, the thesis of Ahmed (1999) proposes a logistic-regression model to determine the probability of a driver choosing to merge, and Toledo et al. (2003) propose an integrated lane-changing model that works for mandatory and discretionary considerations based on a gap-acceptance model.

On the *empirical* side of the literature, many works have focused on how drivers behave at merging points, for instance Chang and Kao (1991), Hounsell et al. (1992), Hidas (2005), Liu and Hyman (2012),and Knoop et al. (2012), Schakel et al. (2012). Of particular interest and relevance to this paper are real-world observational studies that link lane-changing behavior to oscillations between high and low velocities, such as Mauch and Cassidy (2002), Ahn and Cassidy (2007), and Zheng (2014), who find that their data does not support the classical explanation that oscillations arise from car-following, but instead are strongly correlated to lane-changing behavior. However, the oscillations that these papers study are on the target lane, and result from the disturbances that merging creates. In contrast, our results are surprising in that algorithm-driven vehicles on the *blocked* lane may oscillate their velocities.

---

[1] We note that, in some countries, "zipper" merge is recommended in forced-merging scenarios. Nevertheless, these regulations, as well as their enforcement, are inconsistent across different regions and countries, and their application cannot be taken for granted in the grand scheme of things. We do not include cooperative behavior explicitly in our stylized theoretical model. However, driver cooperation is part of our SUMO traffic micro-simulation framework.

In Appendix B we analyze the widely studied NGSIM data-set specifically for oscillations during merging and find that they are quite common, with a big part of them on the blocked lane. We note that, to the best of our knowledge, there are no data-sets or empirical studies on algorithm-assisted driving yet, so the novel cause for traffic oscillations that we uncover here cannot be directly attributed through an observational data-set such as the NGSIM.

On the *theory* side, the merging problem is the topic of early influential papers in Operations Research, such as Weiss and Maradudin (1962), Evans et al. (1964), Hawkes (1968) and McNeil and Smith (1969). These papers build detailed models of merging, and compute quantities of interest such as the merging delay from on-ramp to highway. After a rather long hiatus, this academic community picked up this important topic once again, e.g., Jain and Smith (1997) and Heidemann (2001) from a queueing-theoretic viewpoint, as well as Gregoire et al. (2015), Le et al. (2015) and Como et al. (2016) from a stochastic networks perspective. However, the individual behavior and incentives of drivers are not captured in these works.

Merging is essentially an optimal stopping problem with similarities to the classical "parking problem"; see McQueen and Miller Jr. (1960). The main difference in our stylized model is the existence of an additional decision variable, the velocity, which affects in a non-trivial way both the travel time and the probability of finding a large enough gap to merge. This makes for a richer mathematical structure that allows for new insights such as the existence of velocity oscillations.

The Systems and Control community has viewed vehicular traffic, specifically merging at highway on-ramps, as an optimal control problem. These works are based on either macroscopic, population-level models, e.g., Alessandri et al. (1998), Kotsialos and Papageorgiou (2004), Agarwal et al. (2015), Iordanidou et al. (2015), Pasquale et al. (2015), or on more granular, microscopic models, e.g., Athans (1969), Kachroo and Li (1997), Raravi et al. (2007), Awal et al. (2013), Rios-Torres et al. (2015), Nilsson et al. (2016). While providing a useful benchmark, this viewpoint, again, fails to capture the incentives and global optimization behavior of individual drivers. A decentralized approach, where drivers decide on how to merge based on communication and cooperation has been recently proposed within the framework of Model Predictive Control in Cao et al. (2014) and Cao et al. (2015). We note that there is also a large number of less principled, heuristic approaches to merging in this strand of literature; we refer the interested reader to the survey paper Rios-Torres and Malikopoulos (2017).

The literature in Economics emphasizes the behavior and incentives of rational self-interested drivers, focusing on the resulting traffic equilibria. Notable are the studies on the inefficiencies of vehicular traffic networks with selfish agents, e.g., the Braess paradox; see Wardrop (1952), Vickrey (1969), de Palma et al. (1983), Arnott et al. (1990), Arnott et al. (1991), Correa et al. (2004), Acemoglu et al. (2007) and Acemoglu et al. (2018). Nevertheless, this strand of literature is macroscopic in nature, intended to give high-level insights at equilibrium, and do not focus on the dynamics of traffic and the decisions of individual drivers in the short run.

Tarko et al. (1999) and McCoy and Pesti (2001) have documented *experiments* with "zipper merging" and "late merging" strategies, and have reported on improvements over the uncontrolled random merging benchmark. A similar approach is also taken in Grillo et al. (2008). These works advance our understanding of typical traffic scenarios, but lack an explicit control/optimization dimension, as well as any theoretical support of their findings. Baykal-Gürsoy et al. (2009) and Duret et al. (2010) can be viewed as a step in that direction: the authors introduce theoretical models for predicting the effect of lane-changing using kinematic-wave theory, with a future plan to use these results for better decisions in traffic management.

Finally, given the recent interest in driverless cars, a number of researchers have turned their attention to connected autonomous vehicles and how they interact with manually driven vehicles: Buisson et al. (2018) suggests that autonomous vehicles might harm as well as improve the conditions of on-ramp merging, in the case of a "mixed" traffic flow. Similar conclusions are drawn in Chen et al. (2017), in a more general road configuration. Zheng et al. (2020) demonstrates that flow instability arising in car-following dynamics can be significantly reduced in a "mixed" flow. We complement these recent works by providing theoretical support, in terms of introducing a stylized model and obtaining analytical insights, regarding the phenomena that may occur when an algorithm-assisted driver interacts with other drivers, whose behavior cannot be predicted ex ante. Furthermore, Shalev-Shwartz et al. (2016) and Shalev-Shwartz et al. (2017) propose a Reinforcement Learning-based approach to algorithmic driving. While our approach has several similarities to the latter papers, we focus on (mandatory) lane-changing and provide analytical results and insights via a stylized model.

3. Modeling and analysis of mandatory lane-changing

In this section, we introduce a stylized stochastic model of mandatory lane-changing that is tailored to algorithm-assisted drivers. Through mathematical analysis and extensive numerical experiments, we provide structural results and insights into the dynamics of the merging process.

3.1. Mandatory lane-changing under uncertainty

Consider a two-lane road segment, with traffic moving in the same direction on both lanes. At some point one lane ends, while traffic on the other lane continues to flow freely. Numerous real-life scenarios correspond to this description, such as blockage of the lane due to accident, construction or maintenance works, on-ramp acceleration lanes, and lane drops. We refer to Fig. 1 for a generic illustration of the modeled situation. We denote the two lanes involved as B (Blocked) and F (Free), thus using terminology based on a scenario of accident/maintenance works, but by no means we restrict ourselves to this case. Every vehicle that happens to be on lane B needs to merge to lane F at some place before the blockage point, i.e., we have a scenario of *mandatory lane-changing*.
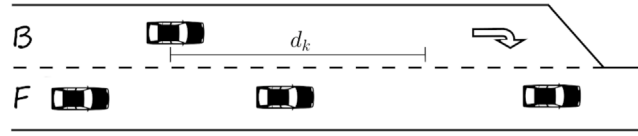
Fig. 1. Illustration of a mandatory lane-changing scenario.

Table 1
Summary of notation.

Notation	Represents
B	Blocked lane
F	Free lane
d_k	Distance between merging positions k and $k+1$
v_k^s, v_k^e	Velocity of the merging car at the start and at end of stage k, respectively
v_H, v_L	Upper and lower bounds, respectively, on the velocity of the merging car
u_k	Binary decision variable of driver's intention to merge
$\bar{v}^F$	Velocity of free lane
Y_k	Bernoulli random variable, denoting whether there is a sufficient gap to allow merging at end of stage $k-1$
$q(\cdot)$	Expected value of Y_k, i.e., merging probability as a function of velocities and distances between merging points
$c(\cdot)$	Expected impact of the relaxation process, in time units, on the merging vehicle
P	Expected penalty, in time units, in case the vehicle reaches the blockage point without having merged
$T_k^B(\cdot)$	Value function of the DP, i.e., optimal expected travel time from stage k, on the blocked lane, to the blockage point

We focus on the behavior of a single vehicle on the blocked lane, whose goal is to merge to the free lane and get through the bottleneck in the shortest possible time. The timing of decision-making is fundamentally discrete, as the blocked-lane driver needs time to prepare for the merging attempt, to adjust velocity if necessary, and to execute safely a merging maneuver, if she so wishes. Potential merging attempts are, therefore, assumed to be sufficiently apart.

Let the starting point of the vehicle be at merging position 0, and the bottleneck be at merging position $N+1$. We denote the distance between positions k and $k+1$ by $d_k, k = 0, \ldots, N$. This discretizes the road segment into cells that we call *stages*, and we define stage k as the road segment between positions k and $k+1$. The exact positions of the gaps on the free lane are not tracked, so at stage $k+1$ the blocked-lane driver might be attempting to merge with the same free-lane driver as at stage k, or it can be a different one (if the flow on the free lane is slow enough). To simplify the notation and analysis, we assume that $d_0 = d_1 = \cdots = d_N = 1$, although this is done without loss of generality.

The vehicle reaches stages 1 to N sequentially. The driver on lane B has two decisions to make at the start of stage $k \in \{0, \ldots, N-1\}$: (a) whether to attempt to merge to lane F at the end of stage k; and (b) whether to accelerate, decelerate, or keep the same speed during stage k, effectively determining the velocity of the vehicle at the end of stage k. Defining the start and the end of each stage is a way of modeling the delay between a decision and the realized outcome of that decision. There is no temporal or spatial difference between the end of stage k and the start of stage $k+1$ for the merging driver.

We use a binary decision variable, u_k, whose value is 1 if the driver's decision is to try merge, and the value 0 otherwise. Let v_k^s and v_k^e denote the velocity of the vehicle at the start and at the end of stage k, respectively. In our benchmark model, we assume that these velocities can only take two values, a high and a low one, v_H and v_L, respectively. (In numerical experiments that, for brevity, we report in Appendix D, we show that our insights hold even when the velocity decision is continuous.)

Three important assumptions underlie our model:

1. The driver on lane B cannot predict with certainty the state of lane F at the end of a given stage, i.e., before she actually reaches that point. In other words, while the *ex ante* start/end of the stages are synonymous to merging opportunities, they may or may not materialize when the blocked-lane driver actually reaches there. This is because the size of the gap depends on the driving behavior of vehicles on the free lane, specifically the vehicles on either side of the gap, which cannot be predicted accurately since there is no communication/coordination between vehicles. In particular, these vehicles may or may not cooperate with the merging driver, to create the necessary space that would be required for merging, which is an especially important factor when traffic density is high.
2. Decisions and actions are not executed instantaneously. Instead, drivers plan a bit ahead to execute an action. In our case, the driver decides to merge at the start of the stage and the decision, if it is to merge, is attempted at the end of the stage. Such a lag, in both information and action, is certainly realistic in human and algorithm-assisted driving, and may be reduced but it is unlikely to go away completely even in fully automated driving.
3. For safety, convenience, and fuel efficiency, a constant acceleration/deceleration rate is applied whenever a change in velocity is mandated. In particular, if varying acceleration/deceleration patterns were allowed, the fastest way to reach the end of a stage with a target velocity could very well be to accelerate and subsequently decelerate within the stage. These trivial, in some sense, oscillations are *not* the focus of our work. In contrast, we uncover that velocity oscillations can arise across stages, even if drivers maintain a constant acceleration/deceleration within each stage.

We also assume that once a vehicle merges to lane F, it never merges back to the lane B. Indeed, when drivers are aware of a blockage, constant lane-changing is rarely observed in practice. As a result, conditional on a successful merging, the expected remaining travel time is a simple function of the (average) velocity of lane F, $\overline{v}^F$. More specifically, the time to traverse stage k on lane F is equal to $1/\overline{v}^F$. We assume that $\overline{v}^F \leq v_L$, i.e., we have a relatively congested free lane, which is typical in mandatory lane-changing scenarios, and the more meaningful regime of analysis for us. In Appendix D we relax this assumption, allowing for $v_L < \overline{v}^F$. In contrast, if the merging vehicle is on lane B, the time that it takes to traverse stage k is equal to $2/(v_k^s + v_k^e)$, i.e., reciprocal to the vehicle's average velocity, as we assume constant acceleration or deceleration.

Let Y_{k+1} be an indicator random variable that takes the value $y_{k+1} = 1$ if a sufficiently large gap (to allow merging) is realized at the end of stage k on the free lane, at the point in time when the blocked lane driver is attempting to merge; and $y_{k+1} = 0$, otherwise. Following the literature, we model the random variable of finding a gap to merge $Y_{k+1} = q(v_k^e, \overline{v}^F, d_k)$, with a (stationary) function $q(\cdot)$ of the velocity of the blocked-lane vehicle, the (average) velocity on the free lane, and the distance between the kth and the $(k+1)st$ merging position. The probability $q(v_k^e, \overline{v}^F, d_k)$ essentially incorporates all uncertainty related to the behavior of the drivers on the free lane, such as whether they are willing to cooperate and create a sufficient gap or not. Even if they do, the gap can be considered by the blocked-lane driver too small to perform a safe merging. We let $q(v_L, \overline{v}^F, 1) = q_L$ and $q(v_H, \overline{v}^F, 1) = q_H$, and we assume that $q_L > q_H > 0$. This is motivated by the fact that, all else being equal, the driver requires less space to merge if the difference between the velocities in the two lanes is small; see Lee (2006), Choudhury et al. (2010) and Toledo et al. (2003) for exact estimates of this effect; and in our case v_L is closer to $\overline{v}^F$, as our main interest is in scenarios where the free lane is relatively dense. We discuss the empirical background and modeling aspects of merging, such as gap-acceptance, in more detail in Appendix C.

A merging vehicle creates a local disturbance in the flow of the target lane; in our case, lane F. This disturbance affects one or more vehicles on the target lane, as well as the merging vehicle itself. This phenomenon, termed *relaxation* in the literature, is well studied and understood empirically; see Zheng et al. (2013) for a recent review. We do not attempt to model the details of the relaxation process as our focus is on the merging vehicle and its objective of minimizing its total travel time. Hence, for simplicity and tractability, we capture the aggregate impact of relaxation on the merging vehicle by a *merging penalty* which, following the literature, we assume to be a (stationary) function $c(v_k^e, \overline{v}^F, d_k)$ of the velocity of the B-lane car and of the F lane, as well as of the distance between the kth and the $(k+1)st$ merging position. We let $c(v_L, \overline{v}^F, 1) = c_L$ and $c(v_H, \overline{v}^F, 1) = c_H$, and we assume that $c_H < c_L$. The reasoning is that, on average, merging requires a larger gap if $v_k^e = v_H$, but conditional on a successful merge event, the relaxation process is smoother overall. In-depth studies of relaxation are limited, but statistical support for the aforementioned claim can be found in van Beinum et al. (2018). A more detailed discussion of relaxation and merging penalties can be found in Appendix C.

If the driver reaches the bottleneck stage without having merged to lane F, she incurs an expected late-merging penalty of P time units, and then exits the bottleneck. (In this late-merging penalty we also include the traveling time from the last merging opportunity, at position N, to the bottleneck, at position $N + 1$.) This penalty models the fact that the vehicle needs to come to a complete stop, and then wait for a large enough gap on lane F in order to bypass the blockage. The regime of interest is when P is reasonably large, as otherwise there may be no incentive to merge. For technical reasons, we assume that the late merging penalty satisfies $P \gg 1 / \min\{v_L, \overline{v}^F\} + c_L$, which is consistent with the intuition that penalty needs to be "large enough".

Finally, as we mentioned already, an algorithm-assisted driver should be viewed as a rational, self-interested decision maker, whose objective is to minimize her (expected) total travel time until she passes the bottleneck, based on the information available to her.[2] In that case, the blocked-lane driver's problem can be formulated as a *Dynamic Programming* (DP) problem, with the following primitives.

State: The state of the system at the start of stage k consists of the lane on which the vehicle is located, F or B, and its current velocity v_k^s, in case the vehicle is still on the B lane.

Decision: The decision at stage k, assuming that the vehicle is still on lane B, consists of two components. The first component represents the driver's intention to merge at the end of stage k, $u_k \in \{0, 1\}$. The second component is the velocity to achieve at the end of stage k, $v_k^e \in \{v_L, v_H\}$. If the vehicle has already merged to lane F at stage k, then there is no decision to be made.

Uncertainty: The uncertainty is on the existence of a sufficient gap on lane F at the end of stage k (equivalently, at merging position $k + 1$), which is captured by the indicator random variable Y_{k+1}.

Dynamics: If the vehicle is on lane B at the end of stage k, and $u_k = y_{k+1} = 1$, then the car merges to lane F at the start of stage $k + 1$. Otherwise, the vehicle remains on lane B and moves at velocity $v_{k+1}^s = v_k^e$.

Objective Function: The objective for the blocked-lane driver is to minimize her expected remaining travel time.

---

[2] Needless to say, human drivers may have different (implicit) objectives in merging scenarios, e.g., safety or ease. However, the total travel time would be the canonical objective of a rational self-interested agent in this setting, as our attempt is to model algorithmic rather than human driving. Furthermore, the expectation of this objective with respect to the available information, i.e., with respect to the driver's beliefs about the surrounding uncertainty, would be the first-order performance metric to consider, consistently with the standard approaches in dynamic optimization under uncertainty.

We denote by $T_k^B(v_k^s)$ the optimal expected remaining travel time, if at the start of stage k the vehicle is on lane B and moves at velocity v_k^s. Similarly, T_k^F represents the optimal expected cost-to-go function if the vehicle is on lane F. We are interested in calculating $T_0^B(v_0^s)$ and, along the way, in obtaining the optimal merging and velocity control policy. This can be done through the Dynamic Programming recursion:

$$T_k^B(v_k^s) = \min_{v_k^e \in \{v_L, v_H\}} \left\{ \frac{2}{v_k^s + v_k^e} + \min\left\{ T_{k+1}^B(v_k^e), q(v_k^e)\left(T_{k+1}^F + c(v_k^e)\right) + (1 - q(v_k^e))T_{k+1}^B(v_k^e) \right\} \right\}, \tag{1}$$

for $k = 0, \ldots, N-1$, with boundary condition:

$$T_N^B(v_N^s) = P, \tag{2}$$

and

$$T_k^F = \frac{N - k + 1}{\bar{v}^F}, \qquad k = 0, \ldots, N. \tag{3}$$

3.1.1. Structure of the optimal policy

Let us denote by (u_k^*, v_k^{*e}) the element of an *optimal* policy at the kth stage. Our initial finding is quite intuitive, and follows directly from the DP recursion in Eq. (1).

Proposition 1. *Consider the mandatory lane-changing problem in Eqs. (1)–(3). If it is optimal not to merge at stage k, i.e., $u_k^* = 0$, then it is optimal to have high velocity during stage k, i.e., $v_k^{*e} = v_H$.*

In other words, the decision pair $(0, v_L)$ can never be part of an optimal merging and velocity control policy. In order to characterize further the structure of optimal policies, we need to introduce following notation:

$$B_k(v_k^s) \equiv \frac{2}{v_k^s + v_L} - \frac{2}{v_k^s + v_H} + q_L\left(T_{k+1}^F + c_L\right) - q_H\left(T_{k+1}^F + c_H\right),$$

and

$$\Delta T_{k+1} \equiv (1 - q_H)T_{k+1}^B(v_H) - (1 - q_L)T_{k+1}^B(v_L).$$

In our formulation, merging can be viewed as an optimal stopping problem, where the driver incurs a lump-sum cost $T_{k+1}^F + c(v_k^e)$ if she merges successfully at the end of stage k, and after which the whole process terminates. In that light, $B_k(v_k^s)$ can be interpreted as the expected loss of choosing v_L over v_H at stage k, while ΔT_{k+1} as the expected continuation benefit of that same decision.

Proposition 2. *Consider the mandatory lane-changing problem in Eqs. (1)–(3). An optimal policy has the following multi-threshold structure:*

1. $(u_k^*, v_k^{*e}) = (0, v_H)$ *if and only if* $T_{k+1}^B(v_H) \leq T_{k+1}^F + c_H$;
2. $(u_k^*, v_k^{*e}) = (1, v_H)$ *if and only if* $T_{k+1}^B(v_H) \geq T_{k+1}^F + c_H$ *and* $\Delta T_{k+1} \leq B_k(v_k^s)$;
3. $(u_k^*, v_k^{*e}) = (1, v_L)$ *if and only if* $T_{k+1}^B(v_H) \geq T_{k+1}^F + c_H$ *and* $\Delta T_{k+1} \geq B_k(v_k^s)$.

Combining the last two cases, we obtain that it is optimal to merge if and only if $T_{k+1}^B(v_H) \geq T_{k+1}^F + c_H$.

Proof. See Appendix A. ∎

The statement of Proposition 2 involves non-strict inequalities, with possibly two optimal solutions on the boundaries. Note that optimal policies have an intuitive structure: if merging is not optimal, something that is associated with moving at high velocity, then the expected remaining travel time on lane F must be larger than that on lane B, even if successful merging was guaranteed, and vice versa. Moreover, if merging is optimal, then the driver should choose v_L over v_H only if the expected continuation benefit of that decision outweighs the expected loss from the merge.

Proposition 3. *Consider the mandatory lane-changing problem in Eqs. (1)–(3). If it is optimal to merge at stage k, i.e., $u_k^* = 1$, then it is optimal to merge at stage $k + 1$, i.e., $u_{k+1}^* = 1$.*

Proof. See Appendix A. ∎

This result can be viewed as the direct extension of the classical result regarding the "parking problem." While such results in optimal stopping problems are usually established by invoking the One-Step Lookahead rule (the "parking problem" is one example), we provide a proof in Appendix A from first principles.

3.2. Traffic oscillations

We now focus on the final region where it is optimal to merge, and also where optimizing driver behavior can lead to velocity oscillations, under certain conditions.

Final Merging Zone: This is the region defined as

$$T_k^B(v_L) \geq T_k^B(v_H) \geq T_{k+1}^F + c_H = \frac{N-k+1}{\bar{v}^F} + c_H. \tag{4}$$

Proposition 3 establishes that there is typically a zone of consecutive stages, reaching the blockage point, where it is optimal to merge. On the other hand, Proposition 2 allows for merging to happen at either velocities, depending on the state of the system and the parameters of the problem. For convenience, let us adopt the shorthand notation

$$\mathbb{E}_{k+1}[L] \equiv q_L \left(T_{k+1}^F + c_L\right) + (1-q_L)T_{k+1}^B(v_L), \tag{5}$$

which denotes the expected remaining travel time at the *end* of stage k, but right before the merging opportunity is revealed (i.e., before the random variable Y_{k+1} is realized), assuming that $v_k^e = v_L$; $\mathbb{E}_{k+1}[H]$ is defined similarly.

The DP recursion can be expressed in the following form:

$$T_k^B\left(v_k^s\right) = \min_{v_k^e \in \{v_L, v_H\}} \left\{ \frac{2}{v_k^s + v_k^e} + \min\left\{\mathbb{E}_{k+1}\left[v_k^e\right], T_{k+1}^B\left(v_k^e\right)\right\} \right\}. \tag{6}$$

The metric of interest here is $\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H]$, i.e., the future loss/benefit from choosing v_L over v_H at stage k. Depending on the exact value of this quantity, the optimal velocity decisions are made as follows:

1. If $\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H] \leq \frac{2}{v_L+v_H} - \frac{1}{v_L}$ and $v_k^s = v_L$, then $v_k^{*e} = v_L$;
2. If $\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H] \geq \frac{2}{v_L+v_H} - \frac{1}{v_L}$ and $v_k^s = v_L$, then $v_k^{*e} = v_H$;
3. If $\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H] \leq \frac{1}{v_H} - \frac{2}{v_L+v_H}$ and $v_k^s = v_H$, then $v_k^{*e} = v_L$;
4. If $\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H] \geq \frac{1}{v_H} - \frac{2}{v_L+v_H}$ and $v_k^s = v_H$, then $v_k^{*e} = v_H$.

To see this, note that if $v_k^s = v_L$ and merging is optimal, Eq. (1) implies that $v_k^{*e} = v_L$, as long as

$$\frac{2}{v_L + v_L} + q_L\left(T_{k+1}^F + c_L\right) + (1-q_L)T_{k+1}^B(v_L) \leq \frac{2}{v_L + v_H} + q_H\left(T_{k+1}^F + c_H\right) + (1-q_H)T_{k+1}^B(v_H).$$

This is precisely the first case above. The other cases can be proved similarly.

It can be verified that

$$\frac{2}{v_L + v_H} - \frac{1}{v_L} \leq \frac{1}{v_H} - \frac{2}{v_L + v_H} \leq 0,$$

since $v_L \leq v_H$. Hence, in the more interesting part of the state space, the merging zone, one can view the optimal policy as having three regions:

Region L: If $\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H] < \frac{2}{v_L+v_H} - \frac{1}{v_L}$, then $v_k^{*e} = v_L$. This is the case when conditions 1 and 3 above are satisfied simultaneously.

Region H: If $\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H] > \frac{1}{v_H} - \frac{2}{v_L+v_H}$, then $v_k^{*e} = v_H$. This is the case when conditions 2 and 4 above are satisfied simultaneously.

Region X: If $\frac{2}{v_L+v_H} - \frac{1}{v_L} \leq \mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H] \leq \frac{1}{v_H} - \frac{2}{v_L+v_H}$, then $v_k^{*e} \neq v_k^{*s}$.

Choosing v_L over v_H always comes at a short-term disadvantage, since it takes longer to traverse the current stage. Thus, the velocity decision depends critically on the future benefit or loss from choosing one over the other. If choosing v_L over v_H is harmful or largely indifferent, then the optimal velocity decision is v_H; this is region H. In contrast, if choosing v_L over v_H is quite beneficial, then the optimal velocity decision is v_L; this is region L.

Proposition 4. *Consider the mandatory lane-changing problem in Eqs. (1)–(3). If the driver enters region L in the merging zone, she stays in region L (i.e., she keeps trying to merge at v_L) until she exits the bottleneck.*

Proof. This result is established as part of the proof of Proposition 5. ∎

Together, Propositions 1 and 4 imply that, in the general case, the optimal policy starts with a region where it is optimal not to merge and move at high velocity, and ends with a region where it is optimal to merge and move at low velocity.

Now, to understand why velocity oscillations may be optimal for an optimizing driver, consider the car on lane B entering stage k with velocity $v_k^s = v_H$. (An identical argument can be made if the car enters stage k at low velocity.) During the next two stages, the driver can act in one of the following three ways (while intending to merge each time, but unable to find sufficient gap to merge):

(a) $H \to H \to H$;

(b) $H \to L \to L$;

(c) $H \to L \to H$.

Sequence (c) is preferable to (a) for the driver if

$$\frac{2}{v_H + v_L} + \mathbb{E}_{k+1}[L] \leq \frac{1}{v_H} + \mathbb{E}_{k+1}[H] \iff \mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H] \leq \frac{1}{v_H} - \frac{2}{v_L + v_H}.$$

Similarly, sequence (c) may be preferable to (b) in the second step because

$$\frac{2}{v_L + v_H} + \mathbb{E}_{k+2}[H] \leq \frac{1}{v_L} + \mathbb{E}_{k+2}[L] \iff \frac{2}{v_L + v_H} - \frac{1}{v_L} \leq \mathbb{E}_{k+2}[L] - \mathbb{E}_{k+2}[H].$$

So (c) is preferable to *both* (a) and (b) when condition X is satisfied. Hence, as long as the quantity $\mathbb{E}_k[L] - \mathbb{E}_k[H]$ stays in region X for consecutive stages, it is optimal for the driver to oscillate between high and low velocity over all the stages in region X: sticking to a low velocity has too high immediate cost, whereas sticking to a high velocity has too high long-term cost, so a rational decision maker attempts to "interpolate" the two extremes, as we explain further below.

3.2.1. The nature of oscillations

Velocity oscillations can arise for different reasons. One obvious reason is the driver fixes her attempted merging points at the (known, estimated, or calculated) gaps on the free lane, targeting optimal velocities at the merging points, and between these merge points minimizes her travel time on the blocked lane. This results in a straightforward bang–bang type behavior, where the driver accelerates and then decelerates to the required velocity between the merging points. These are caused by rational optimizing behavior, but they are not the focus of our work as the phenomenon is not that surprising.

Even if we allow only constant acceleration/deceleration between merging points though, velocity oscillations may arise for two reasons:

1. The driver may unexpectedly see a gap between planned merging points, in which case she may attempt to decelerate and merge on the spot. Again, this phenomenon is not particularly surprising and it is not the focus of our work;
2. Even with constant acceleration/deceleration between decision points, purely for optimization reasons, the optimal velocities at the planned merging points themselves vary, so that we can have patterns such as $v_0^* > v_2^* > v_1^*$ or $v_1^* > v_2^* > v_0^*$, which requires the driver to decelerate and then accelerate, or vice versa, at consecutive stages.

This last phenomenon is counterintuitive and has not been identified in the literature so far.

3.2.2. Convexity and oscillations

At the heart of the existence of region X is the convexity of function $\frac{1}{u}$, representing travel time as a function of velocity, itself a law of physics. Intuitively speaking, the travel time at the average velocity is less than the average of the travel times at high and low velocities:

$$\frac{1}{(v_L + v_H)/2} \leq \frac{1}{2}\left(\frac{1}{v_L} + \frac{1}{v_H}\right), \tag{7}$$

which can be equivalently written as

$$\frac{2}{v_L + v_H} - \frac{1}{v_L} \leq \frac{1}{v_H} - \frac{2}{v_L + v_H}. \tag{8}$$

This "convexity gap", i.e., the difference between the left and right-hand side, is precisely the range of region X.

Now, $\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H]$ is the difference in the expected travel time of arriving at the end of the stage at either velocities, L or H, and can be positive or negative depending on the primitives of the problem. It may also happen to fall (strictly) in the range defined by the region X inequalities. In that case, if the driver happens to be at velocity L at the end of the stage and is unable to merge, then she is better off moving to H and trying to merge at the next stage, as the lower travel time from average speed dominates:

$$\frac{2}{v_L + v_H} - \frac{1}{v_L} < \mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H] \iff \frac{2}{v_L + v_H} + \mathbb{E}_{k+1}[H] < \frac{1}{v_L} + \mathbb{E}_{k+1}[L]; \tag{9}$$

while if she happens to be at H just before the end of stage k, she is better off slowing down to L and try to merge at the end of the next stage:

$$\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H] < \frac{1}{v_H} - \frac{2}{v_L + v_H} \iff \frac{2}{v_L + v_H} + \mathbb{E}_{k+1}[L] < \frac{1}{v_H} + \mathbb{E}_{k+1}[H]. \tag{10}$$

This can happen repeatedly across consecutive stages, as long as the difference of the expected values stays in region X.
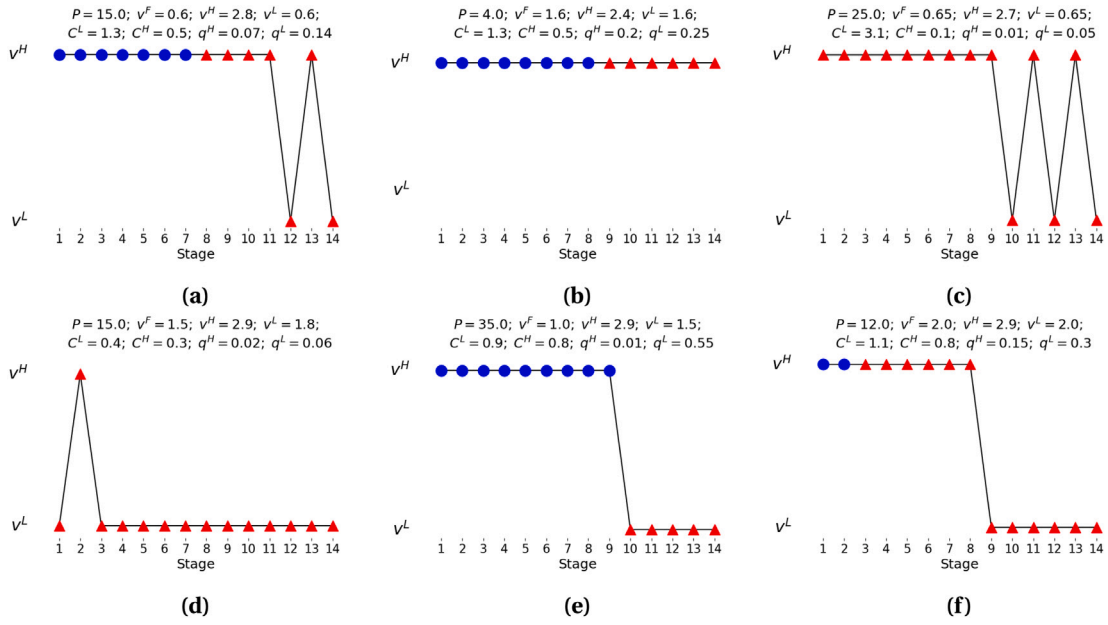
Fig. 2. Examples of the optimal merging and velocity control policy for different choice of parameters P, v_H, v_L, $\overline{v}^F$, q_H, q_L, c_L and c_H, assuming $N = 15$. Optimal merging decisions are represented with the color and the shape of points (blue circle for "not merge", red triangles for "merge"), while the optimal velocities are shown along the y-axis. The corresponding stages are along the x-axis.
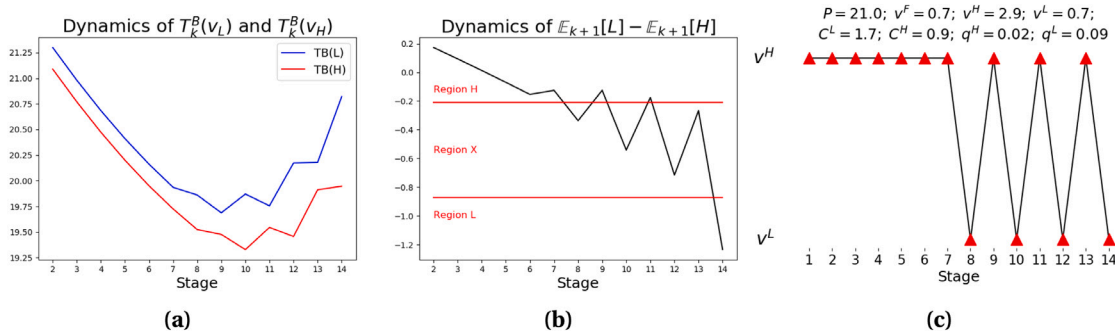

Fig. 3. The dynamics of $\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H]$ lead to an optimal solution that alternates between regions H and X.

3.3. Insight from numerical experiments

To obtain additional insight on whether velocity oscillations occur frequently enough to be of concern, we derive the optimal policy numerically for a broad range of parameter values. Fig. 2 summarizes our findings, illustrating a variety of scenarios that one may encounter. Despite the differences between these cases, a common theme emerges. The optimal policy, in general, seems to have the following three phases, in succession:

$$\text{(no merge, high velocity)} \; \longrightarrow \; \text{(merge, high velocity)} \; \longrightarrow \; \text{(merge, low velocity)}.$$

In Fig. 3 we dive deeper into one of the examples that exhibit the aforementioned velocity oscillations: the left-most figure depicts the expected travel time, at either velocities, for the different stages. Note that after stage 10 the expected travel time increases, due to the fact that incurring the late merging penalty P becomes more and more likely. The center figure shows the evolution, in time, of the quantity $\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H]$, which alternates between regions H and X from stage 7 and onwards. Consistent with our analysis, the optimal policy, which is presented in the right-most figure, exhibits velocity oscillations from stage 7 and onwards.

4. Managing lane-changing of algorithm-assisted drivers

Velocity oscillations while attempting to merge increase the risk of accidents. In this section we show how the traffic manager can avoid global optimization-induced oscillations, by setting appropriate speed limits in the Final Merging Zone.
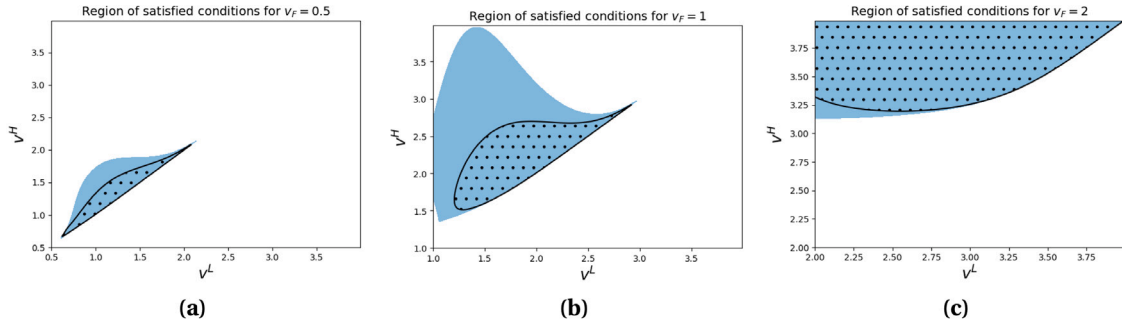
(a)   **(b)**   **(c)**

Fig. 4. A region of parameters v_H, v_L, for which Proposition 5 holds, given reasonable functional forms for q_H, q_L, c_H, c_L, for different values of v_F (in normalized values; they correspond to 2.5 m/s, 5 m/s and 10 m/s). The blue area shows the velocity region that satisfies the first part of proposition, while the dotted area the region that satisfies the second part. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

One way to limit the velocity oscillations is to require that $\mathbb{E}_k[L] - \mathbb{E}_k[H]$ is monotonically non-increasing in k. Then, the dynamics in the Final Merging Zone will be as follows: $H \ldots H X \ldots X L \ldots L$, i.e., with a single region of oscillations. Under stronger conditions, the region of oscillations can be completely avoided.

Proposition 5. *Consider the mandatory lane-changing problem in Eqs. (1)–(3).*

1. *Assume that*

$$\frac{q_L - q_H}{\overline{v}^F} + \frac{2(1 - q_L)}{v_L + v_H} - \frac{1 - q_H}{v_H} + \frac{(1 - q_L)(v_L - v_H)}{v_L(v_L + v_H)} \geq 0,$$

and

$$\frac{q_L - q_H}{\overline{v}^F} - \frac{2(q_L - q_H)}{v_L + v_H} + q_H q_L(c_L - c_H) + \frac{2}{v_L + v_H} - \frac{1}{v_H} + \frac{(v_L - v_H)(1 - q_L)(1 - q_H)}{v_L(v_L + v_H)} \geq 0.$$

Then, $\mathbb{E}_k[L] - \mathbb{E}_k[H]$ is monotonically non-increasing in k, so there can be at most one region of velocity oscillations.

2. *A stronger condition guarantees oscillations never occur. Assume that*

$$\frac{q_L - q_H}{\overline{v}^F} + \frac{2(1 - q_L)}{v_L + v_H} - \frac{1 - q_H}{v_H} + \frac{(1 - q_L)(v_L - v_H)}{v_L(v_L + v_H)} \geq 0,$$

and

$$\frac{q_L - q_H}{\overline{v}^F} - \frac{2(q_L - q_H)}{v_L + v_H} + q_H q_L(c_L - c_H) + \frac{2}{v_L + v_H} - \frac{1}{v_H} + \frac{(v_L - v_H)(1 - q_L)(1 - q_H)}{v_L(v_L + v_H)} \geq \frac{(v_L - v_H)^2}{v_L v_H(v_L + v_H)}.$$

Then, the optimal solution lies in region X for, at most, one stage. Hence, the optimal policy has an intuitive "monotonic" structure:

$$(\text{no merge, high velocity}) \longrightarrow (\text{merge, high velocity}) \longrightarrow (\text{merge, low velocity}),$$

which excludes oscillations.

Proof. See Appendix A. ∎

Interpreting the sufficient conditions on speed limits in Proposition 5 is difficult. However, as there are just two dimensions in our case, v_L and v_H, we can plot the respective regions to (a) check if the conditions have any bite; (b) visualize and determine useful ranges for a policy maker.

In Fig. 4 we provide illustrative examples of the velocity regions that satisfy the aforementioned conditions. First, we set the values for q_L, q_H, c_L, c_H in accordance to functions of $\overline{v}^F, v_H$ and v_L that have appeared in the literature; see Appendix C. Then, for different values of $\overline{v}^F$, we plot the regions (v_H, v_L) where the conditions of Proposition 5 hold. As we can see, in general, the conditions capture reasonably large, implementable regions that, however, need not be convex. Also, the higher the velocity in the free lane, it appears, the larger the area of the parameter space where the optimal solution does include at least one oscillation region.

Let us note that the conditions of Proposition 5 are only sufficient: in Fig. 5 we present indicative examples of optimal policies that avoid oscillations, despite not satisfying the first part of Proposition 5. Specifically, in both examples, the quantity $\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H]$ lies in Region X for exactly one stage.

Pursuing safety may require sacrificing efficiency, leading to larger delays for the drivers. A natural question to ask is the following: if the central planner were to impose safety constraints via speed limits, satisfying the conditions in Proposition 5, what would be the loss in efficiency, measured in aggregate travel time of the drivers? We address this question, numerically, in Fig. 6:
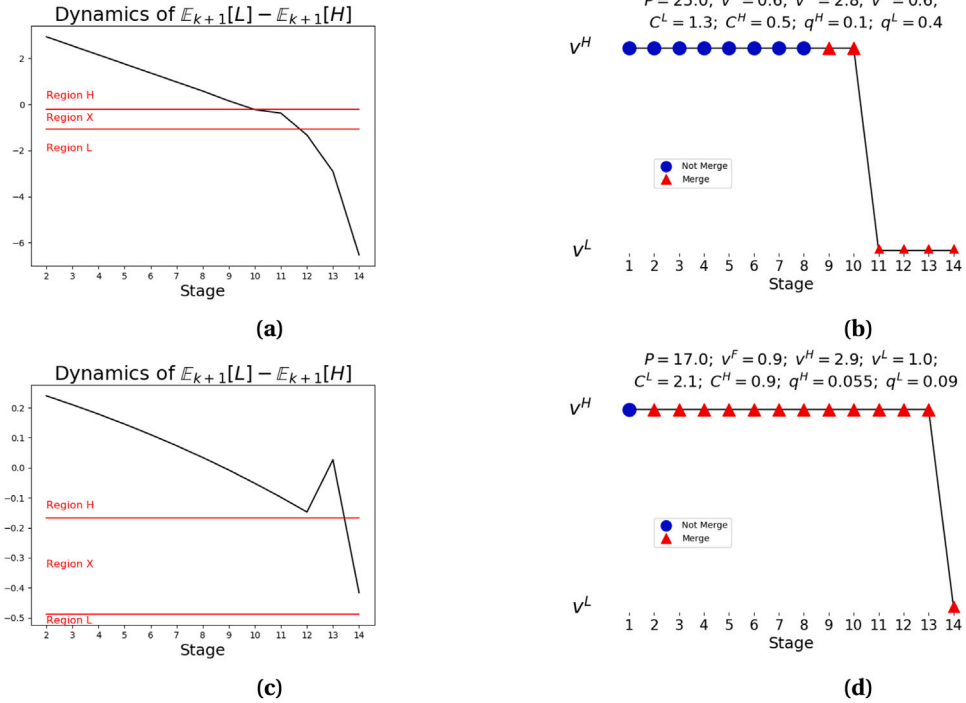
Fig. 5. Two examples of DP solutions that do not oscillate despite failing to satisfy Proposition 5. Top: $\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H]$ behaves smoothly. Bottom: $\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H]$ oscillates, but outside of Region X.
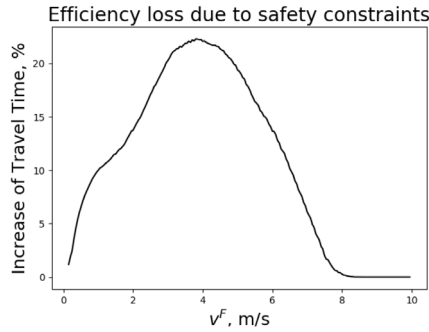

Fig. 6. Additional delay, experienced by merging drivers due to imposing safety constraints, as a function of the velocity on the target lane.

for every such $\overline{v}^F$, we build the "safe" non-oscillating region like the ones in Fig. 4, compute the best policy within the region, and compare it to the best unconstrained policy.

As expected, the loss is near zero when traffic is relatively free, i.e., $\overline{v}^F \geq 8$ m/s. It also goes to zero as the conditions on the free-lane approaches complete jam, i.e., $\overline{v}^F \approx 0$ m/s. The most problematic is the case of intermediately congested traffic, in which case the increase of travel time, experienced by merging drivers, may exceed 20%.

5. Connection to bang–bang control

In this section we consider an extension to our benchmark model on continuous and compact state and action spaces, i.e., the driver may attempt to merge at any point in time, as well as continuously change the vehicle's velocity by accelerating/decelerating. By deriving the Hamilton–Jacobi–Bellman (HJB) equation of the problem, we establish a connections between the optimal policy and bang–bang control. Overall, our treatment in this section is informal, geared more towards providing insight rather than formally proving results.

As before, we consider a car that needs to merge from lane B to lane F in order to get past a blockage point, which lies at position N of lane B. We denote the position and velocity of the merging car by $x \in [0, N]$ and $v \in [v_L, v_H]$ respectively. The

decisions at each point in time are whether to accelerate decelerate or keep the same speed, and whether to merge or not given the opportunity.

As before, we denote the merging decision by $u \in \{0, 1\}$, and the acceleration/deceleration decision by $\alpha \in [-D, A]$, where A and D are the acceleration and deceleration limits of the car, respectively. If the car moves at velocity v on lane B, it finds a large enough gap to merge in lane F with probability $q(v)$. Conditional on merging successfully, it incurs a penalty $c(v)$, which captures the total effect of the relaxation process, and then continues at a constant velocity $\bar{v}^F$ until it gets past the blockage. Functions $q(\cdot)$ and $c(\cdot)$ are assumed to be continuous and differentiable. If the car reaches the bottleneck without having merged, it incurs a (significant) merging penalty P, as it has to come to a complete stop, and then wait for a very large gap on lane F in order to merge.[3]

Let $T^X(x, v)$ be the optimal expected travel time until the car gets past the blockage, assuming that the car is on lane $X \in \{B, F\}$, at position x, and moving at velocity v. Clearly,

$$T^F(x) = \frac{N - x}{\bar{v}^F},$$

since every car on lane F moves at the same velocity $\bar{v}^F$. On the other hand, in the more interesting case where the car is on lane B, we can intuitively express the optimal expected travel time through the following recursion:

$$T^B(x, v) = \min_{\alpha \in A(v)} \left\{ \delta + \min \left\{ q(v + \delta\alpha)\Big(T^F(x + \delta v) + c(v + \delta\alpha)\Big) + \big(1 - q(v + \delta\alpha)\big)T^B(x + \delta v, v + \delta\alpha), T^B(x + \delta v, v + \delta\alpha) \right\} \right\},$$

where $\delta > 0$ is "small" and $A(v) = [0, A]$, if $v = v_L$; $A(v) = [-D, 0]$, if $v = v_H$; and $A(v) = [-D, A]$, otherwise. Furthermore, we have the boundary condition $T^B(N, v) = P$, for all $v \geq 0$.

Let us focus on the region of the state space where it is optimal to merge. (In the region where it is not optimal to merge, it is clear that the car should move at the highest possible velocity.) By taking the Taylor expansion and then letting $\delta \to 0$, we derive the HJB equation of the problem, i.e., function $T^B(x, v)$ is such that

$$\min_{\alpha \in A(v)} \left\{ 1 + q(v)\frac{\partial T^F}{\partial x} v + \big(1 - q(v)\big)\frac{\partial T^B}{\partial v} + \alpha F(x, v) \right\} = 0, \tag{11}$$

where

$$F(x, v) \equiv \frac{\partial q}{\partial v}\Big(T^F(x) + c(v) - T^B(x, v)\Big) + q(v)\frac{\partial c}{\partial v} + \big(1 - q(v)\big)\frac{\partial T^B}{\partial v}.$$

Solving the HJB equation to determine $T^B(x, v)$ is quite challenging analytically. However, the mere structure of the equation provides useful insight into the optimal policy: in order to minimize the left-hand side in Eq. (11), the optimal control needs to be $\alpha \approx -\text{sign}\{F(x, v)\}$. In particular,

$$\alpha = A \cdot 1_{\{v < v_H\}}, \qquad \text{if } F(x, v) < 0;$$

and

$$\alpha = -D \cdot 1_{\{v > v_L\}}, \qquad \text{if } F(x, v) > 0;$$

while any feasible α is optimal, otherwise. This is the essence of the phenomenon that the literature terms bang–bang control: the HJB equation decouples in such a way, so that the product of the control and some state-dependent "drift" term needs to be minimized over a compact set. This implies that the optimal control typically takes values on the boundary of the feasible set (the exact value depends on the sign of the drift term), which allows for, and often implies, oscillatory regions in the optimal policy, as we elaborate below.

First, let us extend the definition in Eq. (5) to continuous space and velocity:

$$G(x, v) \equiv q(v)\big(T^F(x) + c(v)\big) + \big(1 - q(v)\big)T^B(x, v),$$

which denotes the expected remaining travel time at position x and velocity v, right before the merging opportunity is revealed; where $x \in [0, N]$ and $v \in [v_L, v_H]$. Note that

$$\frac{\partial G}{\partial v} = \frac{\partial q}{\partial v}\Big(T^F(x) + c(v) - T^B(x, v)\Big) + q(v)\frac{\partial c}{\partial v} + \big(1 - q(v)\big)\frac{\partial T^B}{\partial v} = F(x, v).$$

In essence, the quantity $\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H]$, on which the analysis of the oscillatory behavior in our benchmark model is largely based, is simply an approximation – in discrete state and action space – to the above partial derivative, and thus to the drift term $F(x, v)$:

$$\frac{\mathbb{E}_{k+1}[H] - \mathbb{E}_{k+1}[L]}{v_H - v_L} \approx \frac{\partial G}{\partial v}\Big|_{(k+1, v_L)} = F(k + 1, v_L).$$

---

[3] This abstract model of merging with continuous state and action spaces has structural similarities to the "rocket-rail car problem" in optimal control theory, where a driver wishes to park a rail car with rocket engines on both ends, with the goal of firing the rockets so as to make the car stop at a precise point in the least amount of time. Importantly, however, there is no merging/optimal stopping aspect to that problem.

The boundaries between regions L, H, and X in the benchmark model are dictated by the differences in traversing a stage at different velocities. In the continuous formulation there is no notion of discrete space (stage) or velocity. Hence, the three regions, in discrete and continuous space and velocity, are as follows:

$$\text{Region L:} \qquad \frac{\mathbb{E}_{k+1}[H] - \mathbb{E}_{k+1}[L]}{v_H - v_L} > \frac{1}{v_L(v_L + v_H)} \iff F(x, v) > 0;$$

$$\text{Region H:} \qquad \frac{\mathbb{E}_{k+1}[H] - \mathbb{E}_{k+1}[L]}{v_H - v_L} < \frac{1}{v_H(v_L + v_H)} \iff F(x, v) < 0;$$

$$\text{Region X:} \qquad \frac{1}{v_H(v_L + v_H)} \leq \frac{\mathbb{E}_{k+1}[H] - \mathbb{E}_{k+1}[L]}{v_H - v_L} \leq \frac{1}{v_L(v_L + v_H)} \iff F(x, v) = 0.$$

The optimal policy in regions L and H is intuitive: if the derivative of the expected remaining travel time, with respect to the velocity, is positive (negative), it is optimal to decelerate (accelerate) as much as possible. In region X on the other hand, every feasible decision is optimal, and oscillations are likely to appear "around" it. For instance, if the solution oscillates between the boundaries of regions L and H, crossing region X but never getting "stuck" there, then we have shifts between maximum acceleration and maximum deceleration, akin to the velocity oscillations that we uncover in our benchmark model.

The above analysis reinforces that regions of maximum acceleration (thus, high velocity) and deceleration (thus, low velocity) are a structural property of the problem, and not the artifact of modeling choices, e.g., discrete vs continuous time and/or state–action spaces. Moreover, it justifies why in the numerical experiments presented in Appendix D, the performance improvement by considering a continuous velocity spectrum – instead of just the highest and the lowest velocity, as our benchmark model does – is practically negligible.

6. Micro-simulations

In this section, we evaluate the performance of the optimal policy derived from the DP formulation of the problem[4] under more realistic traffic conditions, and compare it against benchmark merging policies that are often encountered in practice, e.g., "merge early", "merge late", "merge at a specific (intermediate) point", "merge at a random point" etc. We use the open-source simulation software SUMO (see Lopez et al. (2018)) to construct merging scenarios and test different merging policies. For brevity, we only provide a summary of the results in the main text, and relegate the details of our micro-simulation framework to Appendix C. Through traffic micro-simulations, we investigate the fundamental *flow-density* and *travel time-density* diagrams of the DP policy against the aforementioned benchmarks, under open-boundary conditions.

By and large, the results support our main thesis: *the DP policy provides a principled way to achieve near-optimal performance in a broad variety of traffic conditions, whereas the different benchmark policies may or may not perform well depending on the setting.*

6.1. Flow-density and travel time-density diagrams

Micro-simulations with open-boundary condition allow for understanding the fundamental flow-density and travel time-density diagrams of the DP policy against the various benchmarks, under realistic conditions. Specifically, we compare the average travel time of blocked-lane drivers; the average flow of blocked-lane drivers through the bottleneck; as well as the average travel time and flow of all drivers. In the latter case, effectively, we compare the *bottleneck capacity*, and demonstrate that in most situations the proposed approach does not reduce it. Indeed, a policy that improves traffic conditions on one lane at the expense of other lanes would make little sense.

In Figs. 7 and 8 we plot these performance metrics for the different merging policies, for traffic intensity on the *B*-lane 30% and 60%, respectively, attempting to capture a relatively sparse and a relatively dense scenario. All values were re-scaled relative to the DP policy to make for an easier comparison. Note that the flow metric is well aligned with average travel time in all cases. Whenever the policy exhibits high flow, it has a lower average travel time, which is sensible in long-term simulations.

When it comes to the metrics that concern the blocked-lane drivers only, whether it is flow or travel time, the DP policy exhibits clear superiority over the benchmarks; as theory would predict. Regarding the metrics that concern all drivers, the DP policy is typically the best, or close to it,[5] among the different alternatives. Combined, these findings suggest that the DP policy, on average, makes more efficient use of the available resources.

---

[4] The DP policy relies on estimating the parameters of the model introduced in Section 3, and then solving (numerically) the corresponding DP problem. The parameters of the model are estimated through training data that we obtain from additional micro-simulations, where a blocked-lane driver attempts to merge periodically, without travel-time optimization in mind. We calibrate merging probability, gap-acceptance theory, and other specifications in a similar manner to the NGSIM report; see Choudhury et al. (2006).

[5] The total flow of the DP policy can be slightly lower than that of the "early merge" or "random merge" policies for some densities. This can be attributed to the deliberate slow down of merging cars in the DP policy to increase chances of successful merging. While it improves blocked-lane metrics, it might slightly harm overall performance. Imperfect estimation of the DP parameters can be an alternative cause.
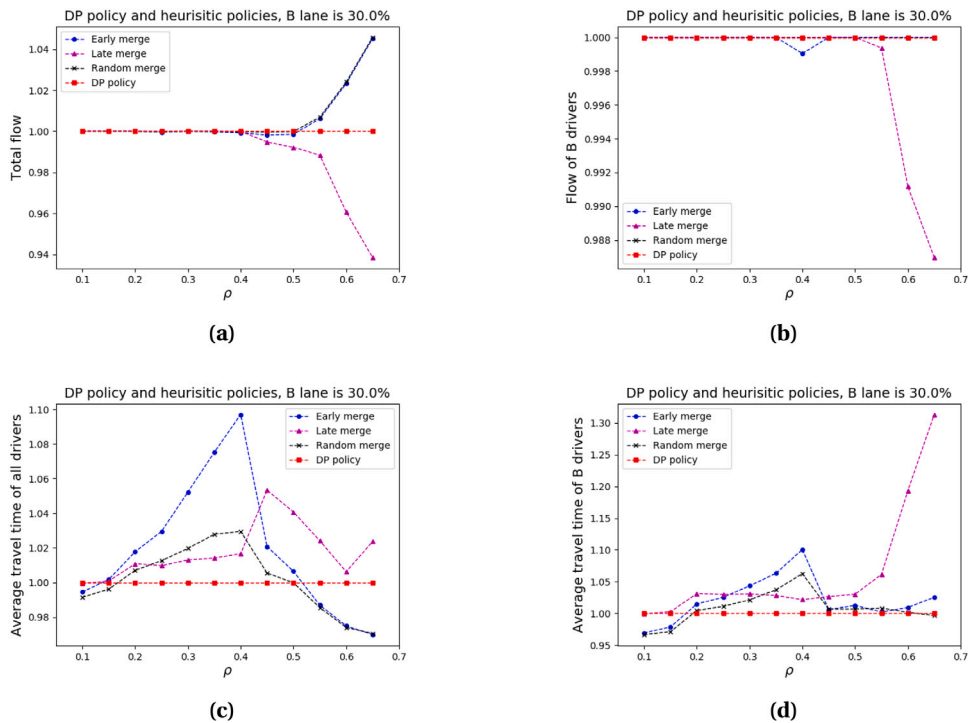
Fig. 7. Comparison of the DP policy to three other popular heuristic policies, for different arrival rates. Traffic on *B*-lane is significant (30% of total arrivals).
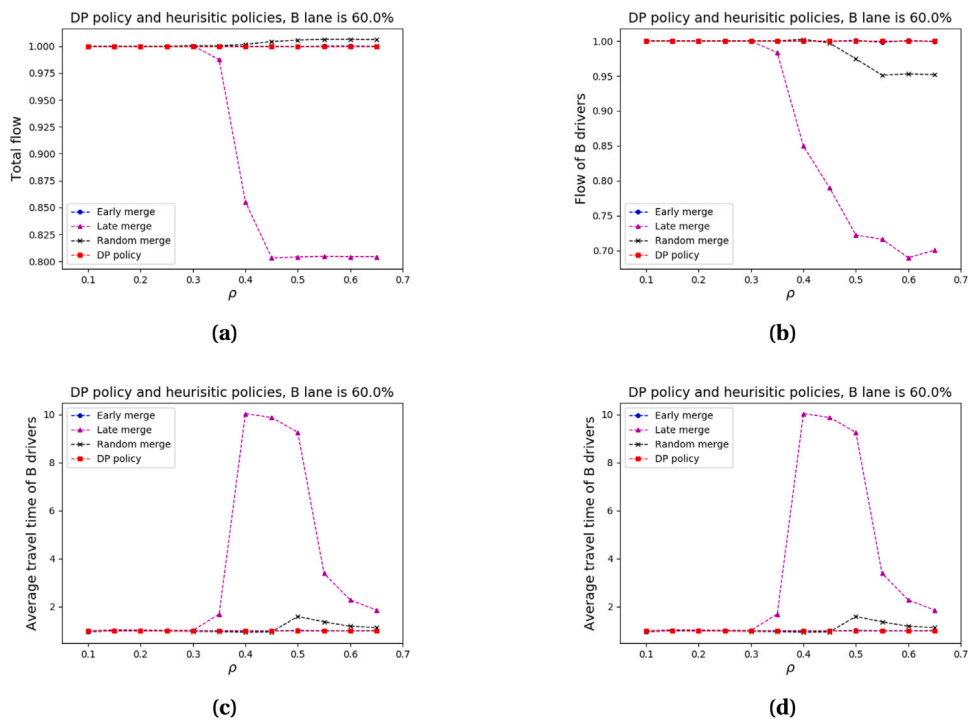

Fig. 8. Comparison of the DP policy to three other popular heuristic policies, for different arrival rates. Traffic on *B*-lane is heavy (60% of total arrivals on *B*-lane).
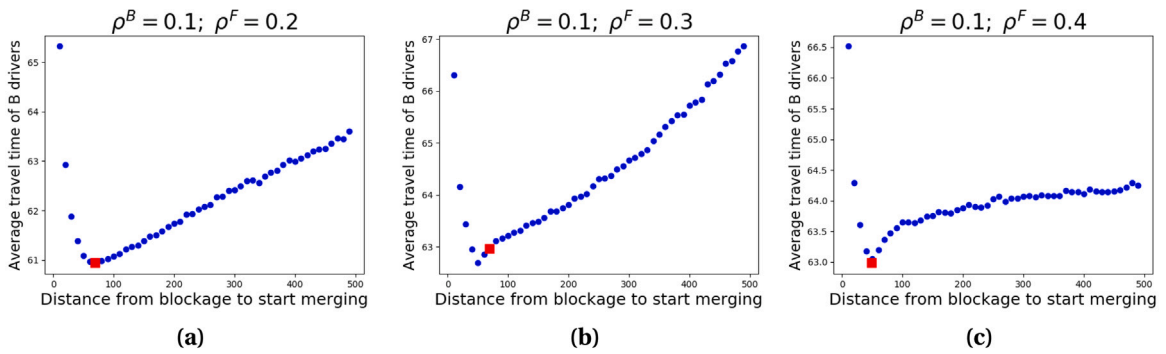
Fig. 9. Comparison of the DP policy to threshold merging policies via traffic micro-simulations. The merging threshold is presented on the *x*-axis, and the total travel time (in seconds) on the *y*-axis. The red square corresponds to the proposed DP policy.

6.2. DP vs. threshold merging policies

Finally, motivated by the desire to understand better any connection between the DP policy and the intuitive benchmark policies "merge early" and "merge late", we analyze the performance of a class of threshold merging policies, which includes the latter as special cases. This class of policies is parameterized by a threshold value, which determines the point on the road segment where the driver begins her attempts to merge, while keeping the velocity as high as possible. In Fig. 9, we present our micro-simulation results regarding the total travel time of these policies as a function of the threshold value, for a relatively dense blocked lane and different densities of the free lane. (Micro-simulations at different densities of the blocked lane exhibit, qualitatively, similar behavior.) Note that the empirical performance of the DP policy is significantly better than policies that merge particularly late or early, and is very close to the (empirically) best one-threshold merging policy.[6] The benefit of our framework is that it provides a principled way to achieve very good performance, irrespective of the traffic conditions, in contrast to the aforementioned class of policies, where the best threshold parameter can only be determined empirically, ex post.

7. Concluding remarks and extensions

Algorithmic-driving and vehicle-to-vehicle communication technologies are reaching maturity, and bring urgency to research on driver behavior, incentives, and control policies for optimal traffic flow. And while decades of research and macroscopic models give good fits to observed phenomena, we still lack a deeper understanding of the underlying causes for fluctuations originated at the driver level, which makes it difficult to devise good intervention and control policies.

To this end, we isolate a simple traffic situation that is reasonably amenable to analysis. Our stylized model provides insights into both the underlying causes as well as potential management policies for algorithm-assisted drivers, i.e., drivers with clearly defined objectives, whose behavioral or cultural tics play less of a role in traffic modeling.

Our modeling framework and DP formulation allows for a characterization of the optimal merging and velocity control policy, resulting in a somewhat counterintuitive finding in some parameter regions: in the presence of uncertainty regarding the future state of the target lane, travel-time optimizing drivers may oscillate between high and low velocities while attempting to merge, a perplexing "irrational" behavior often observed in practice. We validate our theoretical analysis via extensive simulations under real-life scenarios, where we evaluate the macroscopic impact of the DP policy against various merging benchmarks.

This work suggest many potential new research directions. In Appendix D, we provide a few preliminary ideas for extending the work in this paper:

1. First, we consider a direct extension to our benchmark model, where the velocity takes values in a compact set. We provide analytical results and conduct an extensive numerical study, showing that optimal policies with continuous velocity decisions – instead of just the highest and the lowest velocity, as in our benchmark model – are not too different from the ones in our analysis in Section 3.2.

2. Second, we consider an extension to our benchmark where the merging points are not pre-determined but rather specified endogenously. We identify sufficient conditions where the velocity oscillations persist even then.

On the policy and intervention side, apart from setting velocity limits, leveraging technology so as to use prices or set controls on acceleration and distances is expected to lead to significant improvements in traffic control.

---

[6] As expected, an optimal solution to the single-vehicle problem does not always have the best performance, in micro-simulations, among all threshold policies. In some situations, the DP policy recommends merging slightly further from the blockage than the (empirically) best threshold policy, which likely means that either the merging probability is underestimated, or the late merging penalty P is overestimated, or both. The performance gap can also be attributed to unaccounted interactions of the blocked-lane vehicles.

Acknowledgments

Appendix A. Proofs of analytical results

Technical lemmas

We start the appendix by stating certain technical lemmas, which will facilitate the proofs of our main results.

Lemma 1. $T_k^B(v_H) \leq T_k^B(v_L)$, for all $k \in \{1, \ldots, N-1\}$.

Proof. Fix an arbitrary $k \in \{1, \ldots, N-1\}$. The result follows directly from Eq. (1), by noting that the objective function in the optimization problem defining $T_k^B(v_H)$ takes lower values than the objective function in the optimization problem defining $T_k^B(v_L)$, for every feasible solution. ∎

Lemma 2. *If it is optimal to merge at stage k, i.e., $u_k^* = 1$, then*

$$\frac{(v_L - v_H)}{v_L(v_L + v_H)} \leq T_k^B(v_H) - T_k^B(v_L) \leq \frac{(v_L - v_H)}{v_H(v_L + v_H)}.$$

Proof. Let us assume that it is optimal to merge, and that $v_k^{*e} = \tilde{v}$ is the optimal velocity at the end of stage k, if the velocity at the start of the stage is v_H. Eq. (1) implies that

$$T_k^B(v_H) = \frac{2}{v_H + \tilde{v}} + q(\tilde{v})\left(T_{k+1}^F + C(\tilde{v})\right) + \left(1 - q(\tilde{v})\right)T_{k+1}^B(\tilde{v}).$$

Note that,

$$T_k^B(v_L) \leq \frac{2}{v_L + \tilde{v}} + q(\tilde{v})\left(T_{k+1}^F + C(\tilde{v})\right) + \left(1 - q(\tilde{v})\right)T_{k+1}^B(\tilde{v}),$$

since $\tilde{v}$ is a feasible, but not necessarily optimal solution to the optimization problem defining $T_k^B(v_L)$. Together, the two equations imply that

$$T_k^B(v_H) - T_k^B(v_L) \geq \frac{2}{v_H + \tilde{v}} - \frac{2}{v_L + \tilde{v}} = \frac{2(v_L - v_H)}{(v_L + \tilde{v})(v_H + \tilde{v})} \geq \frac{2(v_L - v_H)}{(v_L + v_L)(v_H + v_H)} = \frac{(v_L - v_H)}{v_L(v_L + v_H)},$$

where the right-most inequality uses the fact that the numerator is negative, so that the fraction decreases by substituting v_L for $\tilde{v}$.

This proves the lower bound on $T_k^B(v_H) - T_k^B(v_L)$. The upper bound can be proved by considering that $v_k^{*e} = \hat{v}$ is the optimal velocity at the end of stage k, if the velocity at the start of the stage is v_L, and working similarly. ∎

Proof of Proposition 2. Part (i): Assume that $(u_k^*, v_k^{*e}) = (0, v_H)$ is the optimal solution. This implies that

$$\frac{2}{v_k^s + v_H} + T_{k+1}^B(v_H) \leq \min_{v_k^e \in \{v_L, v_H\}} \left\{ \frac{2}{v_k^s + v_k^e} + q(v_k^e)(T_{k+1}^F + C(v_k^e)) + (1 - q(v_k^e))T_{k+1}^B(v_k^e) \right\}$$

$$= \frac{2}{v_k^s + v_H} + q_H(T_{k+1}^F + c_H) + (1 - q_H)T_{k+1}^B(v_H),$$

where the inequality follows from Eq. (1), for $u_k^* = 0$; and the equality from Proposition 1, which disqualifies $(0, v_L)$ as an optimal solution. Rearranging terms, we have that $T_{k+1}^B(v_H) \leq T_{k+1}^F + c_H$, since $q_H > 0$.

Conversely, assume that $T_{k+1}^B(v_H) \leq T_{k+1}^F + c_H$. We note that

$$\min_{v_k^e \in \{v_L, v_H\}} \left\{ \frac{2}{v_k^s + v_k^e} + q(v_k^e)(T_{k+1}^F + C(v_k^e)) + (1 - q(v_k^e))T_{k+1}^B(v_k^e) \right\}$$

$$\geq \frac{2}{v_k^s + v_H} + \min_{v_k^e \in \{v_L, v_H\}} \left\{ q(v_k^e)T_{k+1}^B(v_H) + (1 - q(v_k^e))T_{k+1}^B(v_H) \right\}$$

$$= \frac{2}{v_k^s + v_H} + T_{k+1}^B(v_H).$$

While the equality is derived through simple algebra, the inequality is based on the following facts: (a) $T_{k+1}^B(v_H) \leq T_{k+1}^F + c_H \leq T_{k+1}^F + c_L$, by assumption and the ordering of merging penalties; (b) Lemma 1, which implies that $T_{k+1}^B(v_H) \leq T_{k+1}^B(v_L)$. Combined with Proposition 1, which precludes $(0, v_L)$ as an optimal solution, the above inequality implies that $(u_k^*, v_k^{*e}) = (0, v_H)$.

Putting the two parts together, we establish the first part of the proposition:

$$(u_k^*, v_k^{*e}) = (0, v_H) \iff T_{k+1}^B(v_H) \leq T_{k+1}^F + c_H.$$

Parts (ii)–(iii): In part (i), we have established that the condition $T_{k+1}^B(v_H) \geq T_{k+1}^F + c_H$ implies that $u_k^* = 1$. Moreover, the condition $\Delta T_{k+1} \leq B_k(v_k^b)$ is equivalent to

$$\frac{2}{v_k^s + v_H} + q_H(T_{k+1}^F + c_H) + (1 - q_H)T_{k+1}^B(v_H) \leq \frac{2}{v_k^s + v_L} + q_L(T_{k+1}^F + c_L) + (1 - q_L)T_{k+1}^B(v_L),$$

if we use the definitions of ΔT_{k+1} and $B_k(v_k^s)$, and rearrange terms. This implies that the high velocity, v_H, is optimal, conditional on being optimal to merge. Combining the two arguments, we have that

$$T_{k+1}^B(v_H) \geq T_{k+1}^F + c_H \text{ and } \Delta T_{k+1} \leq B_k(v_k^s) \iff (u_k^*, v_k^{*e}) = (1, v_H),$$

and

$$T_{k+1}^B(v_H) \geq T_{k+1}^F + c_H \text{ and } \Delta T_{k+1} \geq B_k(v_k^s) \iff (u_k^*, v_k^{*e}) = (1, v_L),$$

which prove the second and third part of the result, respectively.

We emphasize that Proposition 2 does not imply the uniqueness of an optimal solution, i.e., in boundary cases such as $T_{k+1}^B(v_H) = T_{k+1}^F + c_H$ or $\Delta T_{k+1} = B_k(v_k^s)$, there may be more than one optimal solutions. This is the reason that we do not use strict inequalities in the statement or the proof of the result.

Proof of Proposition 3. Assume that it is optimal to merge at stage k, i.e., $u_k^* = 1$. The velocity in the free lane is fixed, $\overline{v}^F$, therefore

$$T_{k+1}^F = \frac{1}{\overline{v}^F} + T_{k+2}^F \implies T_{k+1}^F + c_H = \frac{1}{\overline{v}^F} + T_{k+2}^F + c_H \implies T_{k+1}^B(v_H) \geq \frac{1}{\overline{v}^F} + T_{k+2}^F + c_H, \tag{12}$$

since Proposition 2 implies that $T_{k+1}^B(v_H) \geq T_{k+1}^F + c_H$. Now, Eq (1), with $u_k^* = 1$, implies that

$$T_{k+1}^B(v_H) = \min_{v_{k+1}^e \in \{v_L, v_H\}} \left\{ \frac{2}{v_H + v_{k+1}^e} + \left\{ q(v_{k+1}^e)(T_{k+2}^F + C(v_k^e)) + (1 - q(v_{k+1}^e))T_{k+2}^B(v_{k+1}^e) \right\} \right\}$$

$$\leq \frac{1}{v_H} + q_H(T_{k+2}^F + c_H) + (1 - q_H)T_{k+2}^B(v_H).$$

Combined with Eq. (12), we have that

$$\frac{1}{\overline{v}^F} + T_{k+2}^F + c_H \leq \frac{1}{v_H} + q_H(T_{k+2}^F + c_H) + (1 - q_H)T_{k+2}^B(v_H)$$

$$\iff (1 - q_H)(T_{k+2}^F + c_H) \leq \frac{1}{v_H} - \frac{1}{\overline{v}^F} + (1 - q_H)T_{k+2}^B(v_H)$$

$$\iff T_{k+2}^F + c_H \leq T_{k+2}^B(v_H),$$

since $\overline{v}^F \leq v_H$. This is equivalent to $u_{k+1}^* = 1$, according to Proposition 2.

Proof of Propositions 5 and 4. Our proof strategy relies on establishing sufficient conditions, such that $\Delta E_k \geq 0$, or in other words

$$\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H] \leq \mathbb{E}_k[L] - \mathbb{E}_k[H]$$

over the three regions of the Final Merging Zone, L, X, and H. To that end, we use the fact that

$$\Delta E_k = \frac{q_L - q_H}{\overline{v}^F} + (1 - q_L)\left(T_k^B(v_L) - T_{k+1}^B(v_L)\right) - (1 - q_H)\left(T_k^B(v_H) - T_{k+1}^B(v_H)\right). \tag{13}$$

Region L: If $\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H]$ is in region L, and we are in the merging zone then $v_k^{*e} = v_L$, irrespective of the velocity at the start of stage k. Eq. (13) implies that

$$\Delta E_k = \frac{q_L - q_H}{\overline{v}^F} + (1 - q_L)\left(\frac{1}{v_L} + q_L(T_{k+1}^F + c_L) + (1 - q_L)T_{k+1}^B(v_L) - T_{k+1}^B(v_L)\right)$$

$$- (1 - q_H)\left(\frac{2}{v_H + v_L} + q_L(T_{k+1}^F + c_L) + (1 - q_L)T_{k+1}^B(v_L) - T_{k+1}^B(v_H)\right)$$

$$= \frac{q_L - q_H}{\overline{v}^F} + \frac{1 - q_L}{v_L} - \frac{2(1 - q_H)}{v_L + v_H} - (1 - q_L)q_L T_{k+1}^B(v_L) - (1 - q_H)(1 - q_L)T_{k+1}^B(v_L)$$

$$+ (1 - q_H)T_{k+1}^B(v_H) + (q_H - q_L)q_L\left(T_{k+1}^F + c_L\right)$$

$$= \frac{q_L - q_H}{\overline{v}^F} + \frac{1 - q_L}{v_L} - \frac{2(1 - q_H)}{v_L + v_H} + (1 - q_H)\left(T_{k+1}^B(v_H) - T_{k+1}^B(v_L)\right)$$

$$+ (q_H - q_L)q_L\left(T_{k+1}^F + c_L - T_{k+1}^B(v_L)\right).$$

Note that if it is optimal to merge at v_L, then Eq. (1) implies that

$$\frac{2}{v_k^s + v_L} + q_L(T_{k+1}^F + c_L) + (1 - q_L)T_{k+1}^B(v_L) \leq \frac{2}{v_k^s + v_L} + T_{k+1}^B(v_L).$$

Rearranging terms, we have that $T_{k+1}^F + c_L \leq T_{k+1}^B(v_L)$, which implies that the last term in the expression above is nonnegative (as $q_H - q_L < 0$). Combining this fact with the lower bound on $T_{k+1}^B(v_H) - T_{k+1}^B(v_L)$ provided by Lemma 2, we have the following lower bound for ΔE_k:

$$\Delta E_k \geq \frac{q_L - q_H}{\bar{v}^F} + \frac{1 - q_L}{v_L} - \frac{2(1 - q_H)}{v_L + v_H} + \frac{(1 - q_H)(v_L - v_H)}{v_L(v_L + v_H)} = \frac{q_L - q_H}{\bar{v}^F} - \frac{q_L - q_H}{v_L} \geq 0.$$

Note that if k is in Region L in the merging zone, then $\Delta E_k \geq 0$ as long as $\bar{v}^F \leq v_L$. This also establishes Proposition 4.

Region X: If $\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H]$ is in region X, then it is optimal to merge and $v_k^{*e} \neq v_k^s$. Eq. (13) implies that

$$\begin{aligned}
\Delta E_k &= \frac{q_L - q_H}{\bar{v}^F} + (1 - q_L)\left(\frac{2}{v_H + v_L} + q_H(T_{k+1}^F + c_H) + (1 - q_H)T_{k+1}^B(v_H) - T_{k+1}^B(v_L)\right) \\
&\quad - (1 - q_H)\left(\frac{2}{v_H + v_L} + q_L(T_{k+1}^F + c_L) + (1 - q_L)T_{k+1}^B(v_L) - T_{k+1}^B(v_H)\right) \\
&= \frac{q_L - q_H}{\bar{v}^F} - \frac{2(q_L - q_H)}{v_L + v_H} + (q_H - q_L)T_{k+1}^F + (1 - q_L)q_H c_H - (1 - q_H)q_L c_L \\
&\quad + (1 - q_L)(1 - q_H)\left(T_{k+1}^B(v_H) - T_{k+1}^B(v_L)\right) + (1 - q_H)T_{k+1}^B(v_H) - (1 - q_L)T_{k+1}^B(v_L).
\end{aligned}$$

By adding and subtracting $q_L T_{k+1}^B(v_H)$, we can rewrite the expression as follows:

$$\begin{aligned}
\Delta E_k &= \frac{q_L - q_H}{\bar{v}^F} - \frac{2(q_L - q_H)}{v_L + v_H} + (q_H - q_L)T_{k+1}^F + (1 - q_L)q_H c_H - (1 - q_H)q_L c_L \\
&\quad + (1 - q_L)(1 - q_H)\left(T_{k+1}^B(v_H) - T_{k+1}^B(v_L)\right) + (1 - q_L)T_{k+1}^B(v_H) \\
&\quad - (1 - q_L)T_{k+1}^B(v_L) + (q_L - q_H)T_{k+1}^B(v_H) \\
&= \frac{q_L - q_H}{\bar{v}^F} - \frac{2(q_L - q_H)}{v_L + v_H} + (1 - q_L)q_H c_H - (1 - q_H)q_L c_L \\
&\quad + (1 - q_L)(2 - q_H)\left(T_{k+1}^B(v_H) - T_{k+1}^B(v_L)\right) + (q_L - q_H)\left(T_{k+1}^B(v_H) - T_{k+1}^F\right).
\end{aligned} \tag{14}$$

To find a lower bound on the last term, recall that in region X we have that

$$\frac{2}{v_L + v_H} - \frac{1}{v_L} \leq E_{k+1}[L] - E_{k+1}[H] \leq \frac{1}{v_H} - \frac{2}{v_L + v_H}.$$

Note that

$$\begin{aligned}
&E_{k+1}[L] - E_{k+1}[H] \\
&= (q_L - q_H)T_{k+1}^F + q_L c_L - q_H c_H + (1 - q_L)T_{k+1}^B(v_L) - (1 - q_H)T_{k+1}^B(v_H) \\
&= (q_L - q_H)T_{k+1}^F + q_L c_L - q_H c_H + (1 - q_L)T_{k+1}^B(v_L) - (1 - q_L)T_{k+1}^B(v_H) - (q_L - q_H)T_{k+1}^B(v_H) \\
&= (q_L - q_H)\left(T_{k+1}^F - T_{k+1}^B(v_H)\right) + q_L c_L - q_H c_H - (1 - q_L)\left(T_{k+1}^B(v_H) - T_{k+1}^B(v_L)\right),
\end{aligned}$$

where, again, we add and subtract $q_L T_{k+1}^B(v_H)$ to get the second equality. Using the upper bound for $E_{k+1}[L] - E_{k+1}[H]$ and the equality above, we establish that

$$(q_L - q_H)(T_{k+1}^B(v_H) - T_{k+1}^F) \geq \frac{2}{v_L + v_H} - \frac{1}{v_H} + q_L c_L - q_H c_H - (1 - q_L)\left(T_{k+1}^B(v_H) - T_{k+1}^B(v_L)\right).$$

Combining the above lower bound with Eq. (14), we obtain:

$$\begin{aligned}
\Delta E_k &\geq \frac{q_L - q_H}{\bar{v}^F} - \frac{2(q_L - q_H)}{v_L + v_H} + q_H q_L(c_L - c_H) \\
&\quad + \frac{2}{v_L + v_H} - \frac{1}{v_H} + (1 - q_L)(1 - q_H)\left(T_{k+1}^B(v_H) - T_{k+1}^B(v_L)\right).
\end{aligned}$$

Finally, using the lower bound on $T_{k+1}^B(v_H) - T_{k+1}^B(v_L)$ provided by Lemma 2, we get:

$$\Delta E_k \geq \frac{q_L - q_H}{\bar{v}^F} - \frac{2(q_L - q_H)}{v_L + v_H} + q_H q_L(c_L - c_H) + \frac{2}{v_L + v_H} - \frac{1}{v_H} + \frac{(v_L - v_H)(1 - q_L)(1 - q_H)}{v_L(v_L + v_H)}. \tag{15}$$

Hence, in region X, the right-hand side of the above expression being nonnegative constitutes a sufficient condition for $\Delta E_k \geq 0$;

Region H: If $\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H]$ is in region H, then it is optimal to merge and $v_k^{*e} = v_H$, irrespective of the velocity at the start of stage k. Eq. (13) implies that

$$\Delta E_k = \frac{q_L - q_H}{\bar{v}^F} + (1 - q_L)\left(\frac{2}{v_H + v_L} + q_H(T_{k+1}^F + c_H) + (1 - q_H)T_{k+1}^B(v_H) - T_{k+1}^B(v_L)\right) \tag{16}$$

$$- (1 - q_H)\left(\frac{1}{v_H} + q_H(T_{k+1}^F + c_H) + (1 - q_H)T_{k+1}^B(v_H) - T_{k+1}^B(v_H)\right)$$

$$= \frac{q_L - q_H}{\overline{v}^F} + \frac{2(1 - q_L)}{v_L + v_H} - \frac{1 - q_H}{v_H} + q_H(q_L - q_H)\left(T_{k+1}^B(v_H) - T_{k+1}^F - c_H\right)$$

$$+ (1 - q_L)\left(T_{k+1}^B(v_H) - T_{k+1}^B(v_L)\right)$$

$$\geq \frac{q_L - q_H}{\overline{v}^F} + \frac{2(1 - q_L)}{v_H + v_L} - \frac{1 - q_H}{v_H} + \frac{(1 - q_L)(v_L - v_H)}{v_L(v_L + v_H)},$$

where, to derive the last inequality, we use Proposition 2 and Lemma 2. If the last expression is nonnegative, then $\Delta E_k \geq 0$ in region H.

Summarizing, if both conditions (15) and (16) hold, which is precisely the requirement of Part (i) of Proposition 5, then ΔE_k is nonnegative for every k;

The proof of the second part of Proposition 5 builds on the proof of the first part. Specifically, if the solution lies in region X for at most one stage, then the optimal policy does not exhibit velocity oscillations. A sufficient condition for this to hold is if ΔE_k being greater than or equal to the "length" of region X, for every k. Combining Eq. (15) with the fact that the length of the region X is equal to

$$\frac{(v_L - v_H)^2}{v_L v_H(v_L + v_H)},$$

we derive the following sufficient condition:

$$\frac{q_L - q_H}{\overline{v}^F} - \frac{2(q_L - q_H)}{v_L + v_H} + q_H q_L(c_L - c_H) + \frac{2}{v_L + v_H} - \frac{1}{v_H} + \frac{(v_L - v_H)(1 - q_L)(1 - q_H)}{v_L(v_L + v_H)} \geq \frac{(v_L - v_H)^2}{v_L v_H(v_L + v_H)} \tag{17}$$

Summarizing, if both conditions (16) and (17) hold, which is precisely the requirement of Part (ii) of Proposition 5, then the solution lies in region X at most once, hence the optimal policy exhibits no oscillations.

Appendix B. Analysis of the NGSIM dataset

Description

In this section, we analyze merging instances from the on-ramp in the US Highway 101 Dataset, which is part of the Next Generation Simulation (NGSIM) project of the US DOT.[7] The particular road segment has 5 main lanes with a total length of 2100 ft, as well as an auxiliary lane (lane 6) of 700 ft, starting with an on-ramp and ending with an off-ramp. Video data over three 15-minute intervals (between 7.50 am and 8.35 am) has been transcribed into high-precision snapshots of x and y coordinates and velocities at 100 ms intervals. The data is invaluable for calibration of simulations, and understanding of human driving behavior.

Within this data set, one of the two forced-merging situations in the data is the most relevant to our merging study: vehicles coming from the on-ramp have to merge to the rightmost lane within a specified distance, mimicking a mandatory lane-changing scenario.[8] In our analysis, we focus on two points. Firstly, we provide average velocity statistics on the auxiliary lane versus the rightmost lane (lane 5). That way, we validate our assumption $\overline{v}^F < v_H$. Secondly, we carefully study all forced-merge cases to identify if velocity oscillations happen in practice. To the best of our knowledge, the NGSIM dataset has not been analyzed before with the focus on driver-level oscillations while attempting to merge.

Methodology

To calculate velocity averages, we take an average of velocities of all cars over space and time for the rightmost lane. For the auxiliary lane, we use two different approaches. The first one is the same, taking the average velocity of all cars. However, this might include many cases falling outside the forced merging scenario. It might include cars merging to the auxiliary lane to gain speed advantage or cars merging to the auxiliary lane to reach the off-ramp. Hence, we also provide a second statistic demonstrating the average velocity of 280 identified cases of on-ramp forced merging (prior to the successful merge).

We then turn our attention to identifying velocity oscillations. The main criterion for a merging attempt is a temporal traversal movement towards lane 5, which we identify by checking the difference in the local x-coordinate over a period of time. The x-axis is close to perpendicular (but not perfect) to the direction of the lanes. As a general rule, we classify a traversal movement as a merging attempt, if the vehicle moves towards the target lane at a velocity of at least 5 ft per second, and if such movement is observed for at least 10 consecutive frames (one second).

To identify velocity oscillations, we use the following criteria:

---

[7] https://data.transportation.gov/Automobiles/Next-Generation-Simulation-NGSIM-Vehicle-Trajector/8ect-6jqj.

[8] In total, there are 280 such cases in the NGSIM dataset. One could also consider vehicles merging to the auxiliary lane in order to leave via the off-ramp. However, there were only 97 such cases and they are less homogeneous (for example, multi-lane merges were typical). Hence, we restrict ourselves only to merges *from* the auxiliary lane.

Table 2
Average velocities in US-101 Highway datasets.

Dataset	Rightmost lane	Auxiliary lane, all	Auxiliary lane, on-ramp cars
7.50 am–8.05 am	42.70	54.26	47.90
8.05 am–8.20 am	33.80	50.61	43.17
8.20 am–8.35 am	27.90	46.67	39.15

- We discard cases where the vehicle leaves the auxiliary lane within 4 s, as there is no time for more than one merging attempt;
- We compare the velocity of the merging vehicle to the velocity of the vehicle ahead on the same lane, checking for correlations. This is to rule out the possibility that the merging vehicle simply follows the velocity pattern of the preceding vehicle;
- We track the velocity of the vehicle behind on the target lane as additional confirmation of a merging attempt: in many cases, the vehicle on the target lane reacts by decelerating;
- Our main focus is on the cases where the driver oscillates while merging with the same driver on the target lane, or skips one or more target lane drivers after a failed attempt.

Findings

First, we provide the estimated average velocities (all measured in miles per hour) (see Table 2):

Compared with only on-ramp merging cars, the auxiliary is faster by 12% to 44%. The difference is even more apparent if we take into account all cars on the auxiliary lane. Therefore, the NGSIM dataset supports the assumption regarding the relative velocity in the two lanes of our stylized model.

Further analysis reveals the following phenomena:

1. *Velocity oscillations*, while attempting to merge, take place in about 10% of the forced-merging scenarios. (If we also consider the cases where it takes more than 4 s for the driver to merge, then roughly 20% of the total cases exhibit velocity oscillations.) Usually, these oscillations appear to be relatively small in size (3–7 feet/s) and short in duration (1–2 s). However, there are also occasions where drivers exhibit a consistent and significant velocity drop;
2. *Courtesy yielding*, where a target-lane driver slows down, presumably to create a large enough gap for the merging driver, does occur occasionally but not always;
3. *Blocked merging*, where a target-lane driver not only does not courtesy-yield but instead closes the gap, presumably to prevent merging, also occurs occasionally.

It is, of course, impossible to attribute precise causal factors to the above phenomena from a purely observational dataset, such as the NGSIM one. In Section 3.2.1, we list several reasons due to which oscillations may arise in manual driving.

Since algorithm-assisted driving is based on the global optimization of a clearly defined objective, our theoretical findings highlight the need for 'fair' protocols to be built into these algorithms: if vehicles on the target lane have no incentives to show "courtesy" to merging drivers, then some form of altruistic/cooperative behavior may have to be imposed by a central planner; otherwise, velocity oscillations, detrimental to the overall traffic flow, may arise.

As a specific case study, Fig. 10 illustrates a merging driver from the NGSIM data, and her velocity oscillations while attempting to merge. Shaded (red in color) strips display attempts to merge, determined by significant changes in direction towards the target lane. Note that in the first four merging attempts the velocity oscillates between higher and lower values. We emphasize these are the average velocities at which merging is attempted, and not small variations in velocity which are usually part of the execution of the merging maneuver (as in the first, second, and fourth attempt). The third merging attempt illustrates a blocked merging: the target-lane driver, instead of slowing down to widen the gap, speeds up and matches the merging car's velocity, effectively preventing the merge.

In Fig. 11 we demonstrate four further examples of velocity oscillations of the merging vehicle. A particularly interesting example is vehicle 2122, which changes velocity from low to high, and then back to low, while continuing merging attempts. Vehicle 1789 shows local temporal drops of velocity, which coincide with merging attempts. Vehicles 344 and 831 demonstrate examples of merging attempts at a lower velocity than earlier or later ones.

Appendix C. Numerical experiments

In the stylized, benchmark model presented in Section 3 we made a number of simplifying assumptions for analytical tractability. In this section, we test the proposed DP approach to merging under more realistic conditions. For our simulations we employ SUMO, an open-source framework for discrete-time continuous-space traffic micro-simulations. For a more detailed description of the simulator and its features, see Lopez et al. (2018)

In the analysis of Sections 3 and 4 we assumed that there is a single merging car on the blocked lane; that gaps/merging opportunities are at predetermined positions; that the distance between these positions is the same; and that the velocities of the free-lane cars at different stages are i.i.d random variables. In the simulations that follow we relax all the above assumptions: both the blocked and the free lane have multiple cars; the merging ones follow our DP approach (or different heuristics against which we
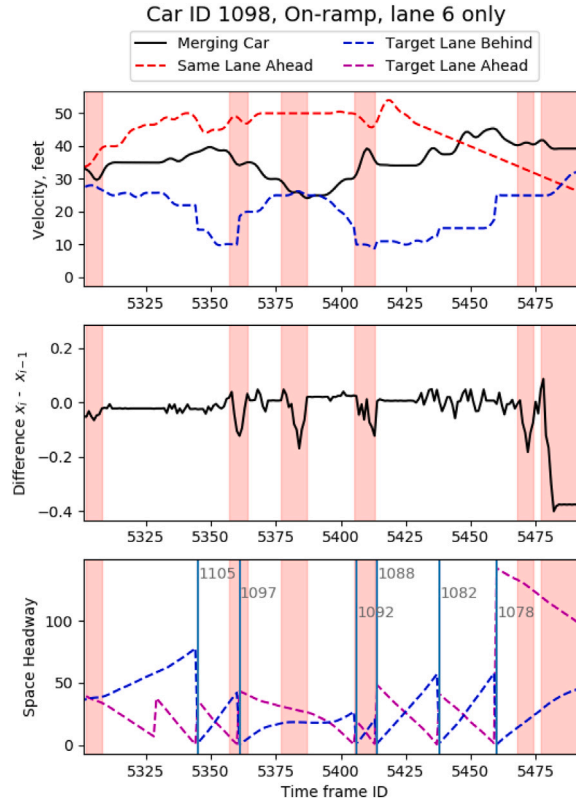
Fig. 10. Driver 1098 of the NGSIM data, showing a temporal drop of velocity correlated with merging attempts. *Top:* Velocity of the merging car, a car behind on the target lane, and a car ahead on the origin lane. *Middle:* Traversal velocity on the lane; negative values signify that the merging car moves towards the target lane, suggesting a merging attempt. *Bottom:* Distance behind and ahead from the merging car to the closest cars on the target lane, signifying the size of the target gap. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

compare our approach), while the ones on the free lane follow a specified set of decision rules that is tuned to mirror characteristics of real vehicular traffic, inducing complicated dynamics about free-lane velocity and merging opportunities.

Before describing the simulation setup, we discuss a numerical approach to gap acceptance, merging probabilities, and merging penalties. They are primarily used in Section 4, although we use merging penalties to adjust DP policy in Section 6.

Gap acceptance and merging probability modeling

While in our theoretical analysis the merging probabilities at different velocities are given exogenously, in order to evaluate the performance of the optimal policy under realistic scenarios, we need to obtain more concrete estimates. The probability of merging depends on the density in the free lane, ρ^F, as a result of the existing gaps in that lane, and the velocities in both lanes. Furthermore, parameter P, which represents the time that the driver has to wait at the blockage point until she finds an appropriate gap, is also a function of the velocity in the free lane. In mathematical terms, in this section we review the existing literature in order to develop functional forms $q(v^B, \overline{v}^F)$ and $P(\overline{v}^F)$, where v^B represents the velocity in the blocked lane, which we can then use, e.g., for blocked-lane velocity optimization.

There is extensive literature of lane-changing behavior based on the *gap-acceptance theory*: a driver who wishes to merge compares the existing gap G at the target lane to a so-called critical gap G_c. If $G > G_c$ then the driver accepts the gap and merges to the target lane, otherwise she stays on the current lane. By adopting the gap-acceptance theory, we define the merging probability as follows:

$$q(v^B, \overline{v}^F) = \mathbb{P}\left[G(\overline{v}^F) > G_c(v^B, \overline{v}^F)\right].$$

Note that this definition is quite flexible, allowing for a broad array of choices for the functional forms of $G_c(v^B, \overline{v}^F)$ and $G(\overline{v}^F)$.

Traffic researchers model the critical gap as made of two components: a lead gap G_c^{lead}, and a lag gap G_c^{lag}. We base our study on the work of Lee (2006), where the following model is developed:

$$G_c^{lead} = \exp(\alpha_0 + \frac{\alpha_1}{1 + \exp(-\max\left\{0, \Delta V^{lead}\right\})} + \alpha_2 \min\left\{0, \Delta V^{lead}\right\} + \frac{\alpha_3 d_{nt}}{1 + \exp(\alpha_4 + \alpha_5 v)} + \alpha_6 v + \epsilon^{lead}), \tag{18}$$
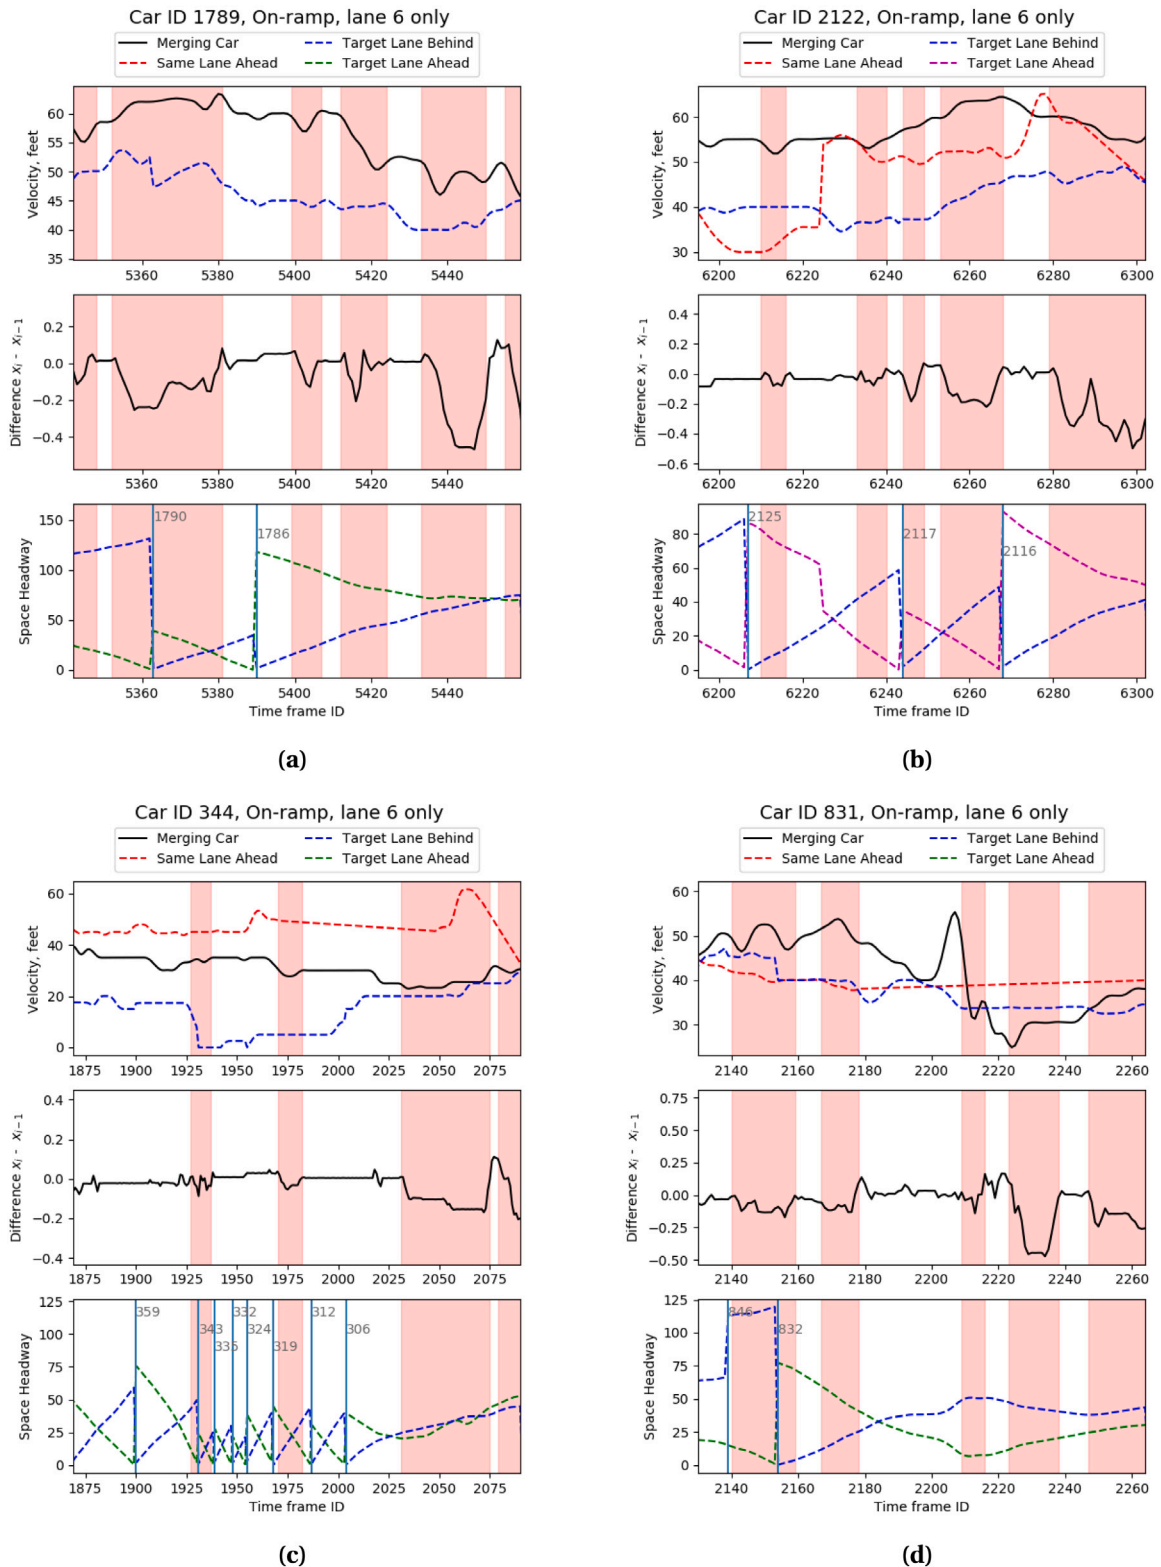
Fig. 11. Velocity oscillations of merging drivers. Each of 4 cases displays 3 subplots. Top: velocity of the merging driver, the velocity of the driver ahead on the same lane, and the driver behind on the target lane. Middle: traversal velocity of the merging car. Bottom: distance to cars behind and ahead (on the target lane), vertical bars show when the driver behind changes, along with a new car ID number. Identified merging attempts are shown in red strips. Time measured in frames, one frame is 0.1 s. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

and

$$G_c^{lag} = \exp(\beta_0 + \beta_1 \max\left\{0, \Delta V^{lag}\right\} + \beta_2 \min\left\{0, \Delta V^{lead}\right\} + \frac{\beta_3 d_{nt}}{1 + \exp(\beta_4 + \beta_5 \nu)} + \beta_6 \nu + \beta_7 \max\left\{0, a^{lag}\right\} + \epsilon^{lag}), \tag{19}$$

where $\Delta V^{lead} = v^{lead} - v^{subj}$, the difference between a lag vehicle and the subject vehicle velocities; d_{nt} is the distance to the end of the lane; ν is a driver-specific aggressiveness parameter; and a^{lag} is the acceleration of the lag vehicle. Lee (2006) also estimates the coefficients α_i and β_i from US highways data, finding them to be significant. A number of more elaborate models for G_c^{lead} and G_c^{lag} have been proposed recently, as well as a different merging models, including courtesy merging (see Choudhury et al. (2010)). We stick here to a simpler model of forced merging.

Since our model ignores the microstructure of the traffic on the free lane, we are not going to divide the whole required gap to lead gap and lag gap. Instead, we will use the pooled gap G_c:

$$G_c = l + G_c^{lead} + G_c^{lag}, \tag{20}$$

where l is the length of the car. (This is due to the fact that in Lee (2006) the lead gap was measured starting from the front bumper, while the lag gap was measured starting from the rear bumper.) To simplify our numerical experiments, we also make the following assumptions:

i. We omit the idiosyncratic aggressiveness term ν, due to the fact that we only have one driver on lane B;
ii. We set $d_{nt} = 0$, suggesting that the critical gap does not depend on the distance from the blockage point in our model;
iii. We omit the factor regarding a^{lag}, as our model does not capture the acceleration decision explicitly;
iv. The random factors ϵ^{lead} and ϵ^{lag} are assumed to be zero-mean normal random variables in the literature, while we set them to zero as an approximation.

Summarizing, G_c is determined by the following equation:

$$\begin{aligned} G_c(v^B, \overline{v}^F) =& 4 + \exp(0.627 + \frac{1.90}{1 + \exp(-\max\left\{0, \overline{v}^F - v^B\right\})} - 0.314\min\left\{0, \overline{v}^F - v^B\right\}) \\ &+ \exp(0.509 + 0.116\max\left\{0, \overline{v}^F - v^B\right\} + 0.034\min\left\{0, \overline{v}^F - v^B\right\}), \end{aligned} \tag{21}$$

where we tentatively use $l = 4$ m for the length of the car.

And while the critical gap is deterministic in our model, the available gap is a random variable. To find an appropriate distribution for G, we first need to estimate the mean density on the free lane, which is typically a function of the velocity, i.e., $d^F(\overline{v}^F)$. Fundamental relationships between velocity and density/headway have been thoroughly investigated in the literature. Following the seminal paper of Greenshields (1934), many authors have offered functional forms for these relationships and have estimated their parameters, e.g., Underwood (1961) and Pipes (1967). We use one of the popular estimated functional forms (see Del Castillo and Benítez (1995)). If we adjust the units of the function to meters and meters per second, then the function takes the following form:

$$d^F(\overline{v}^F) = 7.48\exp(\frac{\overline{v}^F}{8.05}) - l, \tag{22}$$

where, again, we let $l = 4$ m. (We have subtracted the length of the car l as the function determines the headway, i.e., the inter-vehicle distance plus the length of the car.)

It is common to consider headways on a road that are exponentially distributed, e.g., see Miller (1961) and McNeil and Smith (1969). Under our assumption about a fixed velocity in the free lane $\overline{v}^F$, this is equivalent to modeling car arrivals to the studied road section as a Poisson process. Hence, in our numerical experiments we assume that the available gap is an exponential random variable with parameter $1/d^F(\overline{v}^F)$.

Putting everything together, the probability of merging that we use in our numerical experiments is equal to

$$q(v^B, \overline{v}^F) = \exp(-\frac{G_c(v^B, \overline{v}^F)}{d^F(\overline{v}^F)}) \tag{23}$$

where $G_c(v^B, \overline{v}^F)$ and $d^F(\overline{v}^F)$ are given by Eqs. (21) and (22), respectively (see Fig. 12).

Finally, regarding P, the average delay that the driver experiences when stuck at the blockage point, we propose the following model:

$$P(\overline{v}^F) = m + \frac{1}{\overline{v}^F} \cdot \frac{1}{q(0, \overline{v}^F)} + C(0). \tag{24}$$

We interpret the second component as the time that the driver has to wait until she finds an appropriate gap in the free lane, while stuck at the blockage point: $1/\overline{v}^F$ is the average time required to "observe" a new gap, and $1/q(0, \overline{v}^F)$ the average number of attempts needed to find a suitable gap (the mean of corresponding geometric distribution). In our numerical examples we use $m = 2$, as a crude approximation of the time required to actually move from one lane to another. The last component represents the penalty (due to "relaxation") from actual merging to the target lane at 0 velocity. Specific models that can be used for $C(v^B)$ are discussed below.
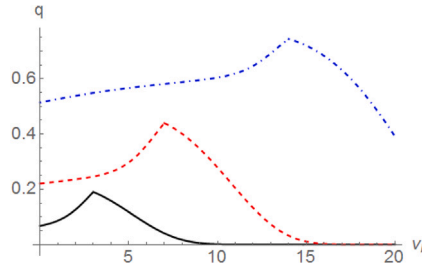
Fig. 12. The function $q(v^B, \overline{v}^F)$ for $\overline{v}^F = 3$ m/s (black solid line), $\overline{v}^F = 7$ m/s (red dashed line), and $\overline{v}^F = 14$ m/s (blue dot-dashed line).

Merging penalty

A merging car creates a disturbance, and both the merging driver and the followers need to adjust velocity and the distance after completing the maneuver. This effect is called "relaxation" and may be an important factor for an optimal merging decision.

We first show important qualitative characterization of the merging penalty, that is, how does it changes with increase of the velocity of merging car.

Let us denote the merging penalty to the driver that merges into gap g at the velocity v^B into the flow with velocity $\overline{v}^F$ as $C(g, v^B, \overline{v}^F)$. As before, we count only the pooled gap, rather than discriminating between lead gap and lag gap. Then, the (expected) merging penalty that we use in our DP can essentially be written as follows:

$$C(v^B) = \mathbb{E}[C(g, v^B, \overline{v}^F)|g > G_c(v^B, \overline{v}^F)]. \tag{25}$$

To show analytically the required characterization, we take several assumptions:

i. Inter-car distances, and therefore, available gaps g are distributed according to the exponential distribution with parameter λ;
ii. Velocities v^B and $\overline{v}^F$ affect the expected merging penalty only through critical gap $G_c(v^B, \overline{v}^F)$, and therefore the merging penalty $C(g, v^B, \overline{v}^F) = C(g)$ depends only on g;
iii. For $g \geq 0$ the function $C(g)$ maps to the non-negative values and is strictly decreasing. We show results a for negative-exponential function $e^{\gamma - \delta g}$, for parameters $\gamma \geq 0, \delta > 0$. We briefly discuss below also an alternative model with function $e^{\gamma - \delta(g - G_c(v^B, \overline{v}^F))}$. These two functions leads to clean analytical results. Qualitative results for reciprocal function $1/(\gamma + \delta g)$ are similar, although less tractable.

We start from merging penalty function $C(g) = e^{\gamma - \delta g}$. Expected merging penalty, conditional on a merging event, can be written then as (we omit variables of critical gap for brevity):

$$C(v^B) = \mathbb{E}[C(g, v^B, \overline{v}^F)|g > G_c] = \frac{1}{1 - F(G_c)} \int_{G_c}^{\infty} f(g)C(g) = \frac{1}{e^{-\lambda G_c}} \int_{G_c}^{\infty} \lambda e^{\gamma - (\lambda + \delta)g} = \frac{\lambda e^{\gamma - \delta G_c}}{\lambda + \delta}. \tag{26}$$

Since it is known (in particular, it follows from the models we discussed above) that $G_c(v_H, \overline{v}^F) > G_c(v_L, \overline{v}^F)$, for some velocities of the blocked lane driver such that $v_H > v_L$, and hence it follows that $C(v_H) < C(v_L)$. Note that this qualitative result remains unchanged, if we would consider a similar functional form for $C(g, v^B, \overline{v}^F)$, but having the penalty depending on the difference between the actual gap and some constant minimally required gap $g_0 < G_c(v^B, \overline{v}^F)$. Indeed, this leads to the result:

$$C(v^B) = \frac{\lambda e^{\gamma + \delta g_0 - \delta G_c}}{\lambda + \delta}, \tag{27}$$

which is the same model with a higher constant.

We consider another special case of the model, when the penalty depends on the difference between the actual gap and the critical gap, i.e. $C(g) = e^{\gamma - \delta(g - G_c)}$. It is straightforward to verify, that in this case the expected penalty would be:

$$C(v^B) = \frac{\lambda e^{\gamma}}{\lambda + \delta},$$

which is independent of v^B. We can conclude, that under reasonable modeling assumptions, ex-post merging penalty in the optimal control and velocity problem is such that $C(v_H) \leq C(v_L)$ for $v_H > v_L$.

We also can use the same model for obtaining the numerical values of $C(v^B)$. As can be seen from relaxation studies, at a high level the merging penalty is the time, lost due to "stretching" leading distance after the lane change to a certain value, normal the for driving conditions on the free lane. Approximately it can be calculated as $(d^F(\overline{v}^F) - g/2)/\overline{v}^F$, that is the time to cover the difference between the gap and the normal for that velocity inter-car distance, at velocity $\overline{v}^F$ (in case $d^F(\overline{v}^F) > g/2$, and zero otherwise). Note that we use only half of gap g, since the merging driver takes into account the lead gap after the merge, and we assume that this is simply the half of the total gap (although it can be not exactly the case in reality).

Table 3
SUMO car-following parameters.

Parameter	Value	Description
tau	1.5	Driver's desired minimum time headway, s
accel	3.0	Maximum acceleration of the vehicle, m/s^2
decel	9.0	Maximum deceleration of the vehicle, m/s^2
length	4.0	The vehicle's length, m
minGap	1.0	Minimum gap between cars in jam, m
maxSpeed	10.0	The vehicle's maximum velocity, m/s
delta	4	Acceleration exponent, IDM parameter
stepping	0.25	The internal step length when computing follow speed, IDM parameter, s

We can use this (linear) model to fit our exponential model $C(g) = e^{\gamma - \delta g}$. To estimate unknown parameters $\gamma \delta$, we need to take two pairs $g, C(g)$ from described above linear model. The estimates will depend on what two points we take; we use here $g = 0$ and $g = 0.5 d^F(\overline{v}^F)$. It is straightforward to see that this leads to the following estimates:

$$\gamma = \ln \frac{d^F(\overline{v}^F)}{\overline{v}^F}; \qquad \delta = \frac{2 \ln 4/3}{d^F(\overline{v}^F)}.$$

Once we obtained the estimates for fixed $\overline{v}^F$, we can calculate ex-post late-merging penalty $C(v^B)$ according to (26). Note that critical gap $G_c(v^B, \overline{v}^F)$ is needed to calculate merging penalty; any appropriate model can be used for it, we use the model discussed earlier (see (21)).

Micro-simulation setup

SUMO is an open-source software providing the framework for discrete time, continuous space traffic micro-simulations. To simulate a blocked-lane scenario, we construct a simple network with a two-lane stretch of the highway. The total length of the road is 700 m, and cars can start merging after initial 200 m of the road. At the end of the stretch, the road is narrowing down to a single lane, and cars can exit the study area only if they merge to the free lane in advance. To model car-following behavior, we use the Intelligent Driver Model (see Treiber et al. (2000)), one of the readily available models in SUMO. The simulation time step is set to 1 s. We list key parameters used in our micro-simulations in Table 3.

We do not use most of the lane-changing behavior available in SUMO. For consistency with our theoretical model, cars are only allowed to merge from the blocked lane to the free lane. Technically, this is achieved by controlling the "lcStrategic" parameter of the lane-changing model. If the car is not expected to change lane, e.g., for late merging policy when the distance to the blockage is large, this parameter is set to zero. Otherwise, it is set to a large number, resulting in repeated attempts to merge to the free lane. All other lane change types, such as getting the speed advantage of a faster lane, or preference to move on the right when traffic is light, are turned off. Collaborative behavior is turned on, so the drivers on the free lane sometimes decide to slow down and create an appropriate gap for the merging car. For further details about the lane-changing model in SUMO, see Erdmann (2014).

Given that closed-boundary conditions do not have a straightforward implementation in SUMO and are, generally speaking, of lesser practical interest, we limit ourselves to open-boundary conditions. If the vehicle reaches the end of the road segment while driving on the free lane, it exits the system. On the other hand, vehicle arrivals are randomly generated on both lanes according to pre-set arrival rates. The simulation ends after 10 000 steps or approximately 3 h of simulated time. After the termination, we calculate all statistics of interest.

Once we have fixed all the parameters of the micro-simulation, we need to evaluate the parameters of the model in Section 3.1, so that we can obtain the optimal solution to the corresponding DP and introduce it as a new set of controls in the simulation framework. We use additional simulations to collect the data and estimate these parameters, similar to what we would be doing if we were collecting real data. More specifically, the estimation method that we use is to repeat the simulations several times, and then take the sample averages of the quantities of interest. Below we provide additional details on the procedure for each DP parameter.

Velocity on the free lane, $\overline{v}^F$. We run simulations repeatedly (without any cars on the blocked lane), keeping track of the velocity of each car at all times. Averaging over the time horizon, the different cars, and the various simulation runs, we obtain an estimate for $\overline{v}^F$;

Velocities v_H and v_L. As discussed earlier, v_L and v_H are parameters that can be chosen by a blocked-lane driver, or imposed by the traffic regulator. The insights from Section 4 suggest that we choose $v_H = v^{max}$ and $v_L \approx \overline{v}^F$, as one of the choices resulting in no oscillations;

Merging probabilities q_H and q_L. To estimate a merging probability, we run repeated simulations with a single driver on the blocked lane, moving at the corresponding velocity and constantly trying to merge. Counting the number of attempts per every successful lane-change (at a given velocity), we obtain an estimate for the merging probability.

Late merging penalty P. While in our theoretical model in Section 3.1 there is only a single car on the blocked lane, in the simulations, the late merging penalty also accounts for the adverse effect of congestion on the blocked lane near the blockage point. To estimate its value, we run repeated simulations with the target density and arrival rates and evaluate the time the car spends
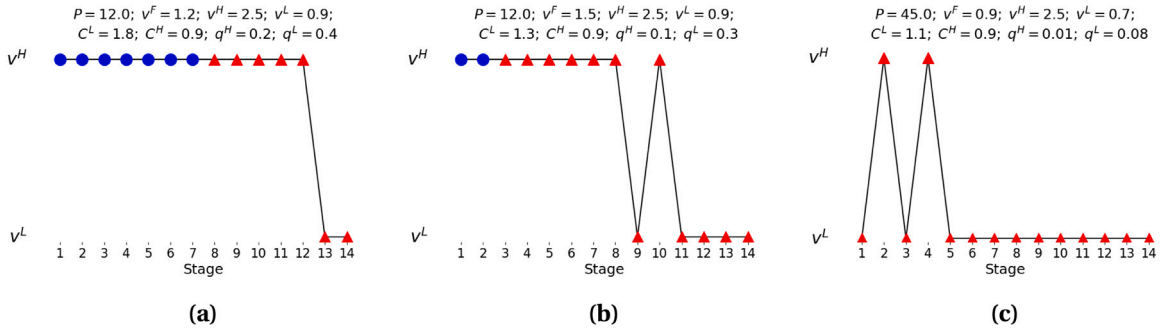
Fig. 13. Examples of the optimal solution for different choice of parameters P, v_H, v_L, $\overline{v}^F$, q_H, q_L, C_H and C_L assuming $N = 15$. In all examples $v_L < \overline{v}^F < v_H$ holds.

in congestion (i.e. with very low velocity) near the blockage (defined as the last 100 m of the road stretch). We denote this as P_0. We then add the analytical component, representing travel time lost due to "relaxation" (we define function $C(\cdot)$ rigorously in the previous section of this Appendix). The final form of the late merging penalty is:

$$P = P_0 + C(0)$$

Merging penalties c_L and c_H: we use the theoretical model, described in the previous section of this Appendix. As a distance function $d^F(\overline{v}^F)$ for free lane we use the one described by (22), and (21) as the model for critical gap $G_c(v^B, \overline{v}^F)$.

Apart from the aforementioned parameters, we also need to fix d, which determines the distance between subsequent attempts to change lane. Note that there is a discrepancy here between the structure of our theoretical model and the simulation. In the latter, a driver can check for merging opportunities once per simulation step. The distance between subsequent attempts clearly depends on the driver's velocity, so it is not constant. On the opposite, there is a known distance between subsequent attempts in our theoretical model. Therefore, the best we can do is to try to solve the single-car discrete-velocity DP for different constant choices of d, which would roughly correspond to the distance the blocked-lane driver covers between simulation steps. The resulting d should not be too small, as this would result in too frequent lane change attempts in the DP policy and a bias towards late merging, but also should not be too large, which would result in too few merging opportunities. In Section 6 we only report on the case where $d = 10$ m, although values ranging from 10 m to 25 m demonstrate similar results.

Appendix D. Extensions

In this appendix we extend our benchmark model in three ways, so as to test the robustness of the derived insights.

  i. First, we relax certain assumptions regarding velocity restrictions, which we make throughout the paper for analytical tractability, by allowing for the low velocity, or both the low and the high velocity, to be less than the average velocity on the free lane. We discuss how this affects our theoretical results, and we demonstrate numerically that in the latter case oscillations may still arise.
 ii. Second, we consider a direct extension to our benchmark model, where the velocity takes values in a compact set. We conduct an extensive numerical study, showing that the optimal policies with continuous velocity decisions – instead of just the highest and the lowest velocity, as in our benchmark model – are not too different from the ones in our analysis in Section 3.2.
iii. Finally, we consider an extension to our benchmark where the merging points are not predetermined but rather specified endogenously. We identify sufficient conditions such that velocity oscillations persist even then.

Relaxed velocity assumptions

In the benchmark model of the paper, we derive theoretical results and insights under the assumption that $\overline{v}^F \leq v_L < v^H$. This case is the most interesting for us, as it represents situations with few drivers on lane B and heavier traffic on lane F. However, one could argue that the driver could also choose a velocity that is lower than $\overline{v}^F$, leading to the case $v_L < \overline{v}^F < v^H$. How would that affect our findings?

First of all, we should be careful about relative velocities. Our assumption about merging probabilities $q_L > q_H$ is based on gap-acceptance theory, and on the assumption that $|v_L - \overline{v}^F| < |v_H - \overline{v}^F|$. While it is automatically satisfied if both v_H and v_L are higher than $\overline{v}^F$, it is not necessary the case if we relax this.

So, let us assume that v_L is still closer to $\overline{v}^F$, so that $q_L > q_H$. In that case, it is straightforward to verify that Propositions 1–3 still hold, but Propositions 4 and 5 do not. Simply put, the merging decision still has a threshold form (if optimal to merge now, then optimal to merge later); oscillations might still exist as Proposition 2 allows them; however, little can be said about sufficient conditions preventing oscillations.
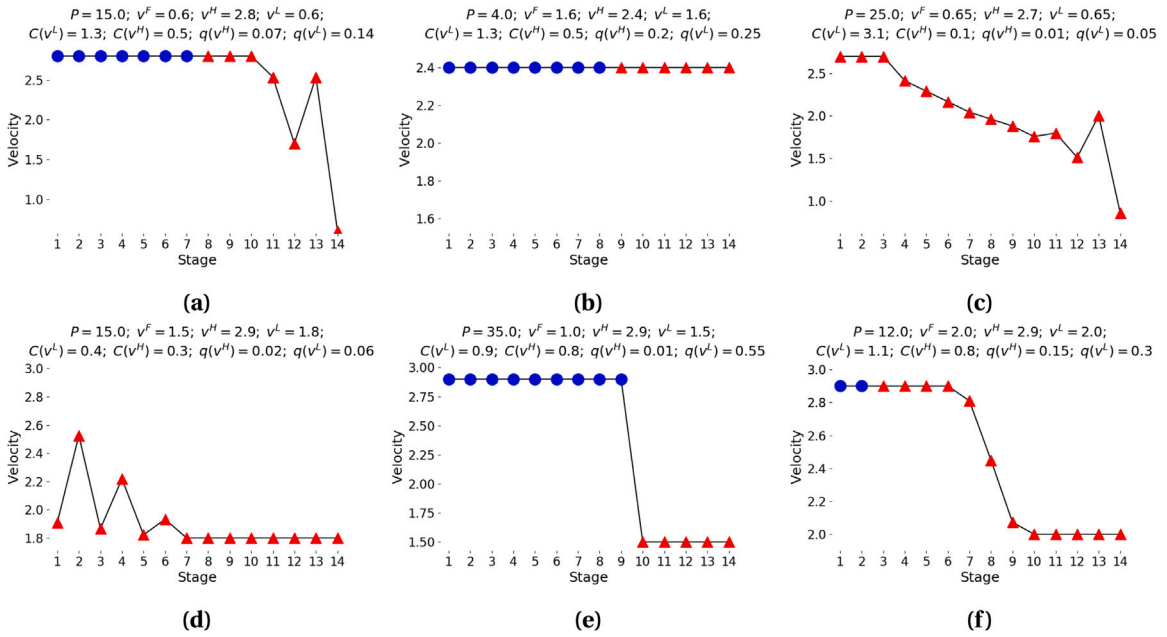
Fig. 14. Examples of the optimal solution for different choice of parameters P, v_H, v_L, $\overline{v}^F$, $q(v_H)$, $q(v_L)$, $C(v_L)$ and $C(v_H)$ assuming $N = 15$. The velocity space is discretized by dividing the interval evenly into 50 discrete levels.

Fig. 13 shows three examples with optimal solutions of DP. They clearly demonstrate that qualitatively there is no much difference. The solution might be non-oscillating in velocity or exhibiting oscillations similar to what could be seen for a model without relaxation of the velocity restrictions. No new types of optimal solutions were found.

Finally, what if $|v_L - \overline{v}^F| > |v_H - \overline{v}^F|$, resulting in $q_H > q_L$? This case may naturally arise if we assume that *both* low and high velocities are lower than $\overline{v}^F$; for example, if the blocked lane is highly congested, or the conditions on the blocked lane do not allow safe acceleration above $\overline{v}^F$. A careful examination of Eq. (1) would reveal that in this case there are no tradeoffs. A decision $\left(u_k^*, v_k^{*e}\right) = (1, v_H)$ dominates all others, and therefore, the driver should always attempt to merge, using only high velocity.

Continuous velocity

Consider a model that is the same as our benchmark one, except for the fact that the velocity is a continuous decision variable, with values in $[v_L, v_H]$. The DP recursion in this case takes the following form:

$$T_k^B\left(v_k^s\right) = \min_{v_L \le v_k^e \le v_H} \left\{ \frac{2}{v_k^s + v_k^e} + \min\left\{ \underbrace{q(v_k^e)\left(T_{k+1}^F + C(v_k^e)\right) + (1 - q(v_k^e))T_{k+1}^B\left(v_k^e\right)}_{u_k=1}, \underbrace{T_{k+1}^B\left(v_k^e\right)}_{u_k=0} \right\} \right\}. \tag{28}$$

The propositions below extend the discrete-velocity results to the case with continuous velocities, and for brevity, we omit their proofs, as they are very similar to the proofs of their counterparts.

Proposition 6. *Consider the mandatory lane-changing problem in Eq.* (28). *If it is optimal not to merge at stage k, i.e., $u_k^* = 0$, then it is optimal to have the highest velocity during stage k, i.e., $v_k^{*e} = v_H$.*

Proposition 7. *Consider the mandatory lane-changing problem in Eq.* (28). *If it is optimal to merge at stage k, i.e., $u_k^* = 1$, then it is optimal to merge at any later stage, i.e., $u_l^* = 1$, $l = k + 1, \ldots, N$.*

Proceeding analytically with Eq. (28) is challenging, as the first-order condition for minimizing the expected travel time requires solving a fourth-order polynomial in v_k^e. Thus, we resort to a numerical investigation, by approximating the compact velocity space by a large number (50) of discrete velocity levels. Also, for the remainder of the section, we adopt a linear merging probability model: $q\left(v^B\right) = \alpha - \beta v^B$, where $\alpha, \beta > 0$ and $\beta v_H < \alpha \le 1$; and a negative-exponential function for the merging penalty: $c\left(v^B\right) = e^{\gamma - \delta v^B}$, where $\delta > 0$. We refer the reader to Appendix C for details on selecting these functional forms.

Fig. 14 presents examples of (approximately) optimal policies in the continuous-velocity case. Note that we use the same parameter sets as in Fig. 2, so that the two scenarios are directly comparable.

A first observation is that, overall, the optimal policies in the continuous-velocity scenario do not differ much from those in the discrete-velocity benchmark. In fact, in some cases, the optimal policies are exactly the same, while in the rest, the difference is

typically confined in a relatively narrow window during which the driver needs to slow down. To investigate this further, we report on the gap between the optimal expected travel times in the two scenarios in the table below. (Again, in the continuous-velocity case, the corresponding DP is solved numerically by discretizing the velocity space in 50 levels, and the linear probability model is adopted.) The first six rows correspond to the parameter sets in Figs. 2 and 14, and three additional cases are considered.

$\bar{v}^F$	v_L	v_H	$q(v_L)$	$q(v_H)$	P	$C(v_L)$	$C(v_H)$	Gap
0.6	0.6	2.8	0.14	0.07	15	1.3	0.5	0.331%
1.6	1.6	2.4	0.25	0.2	4	1.3	0.5	0%
0.65	0.65	2.7	0.05	0.01	25	3.1	0.1	1.406%
1.5	1.8	2.9	0.06	0.02	15	0.4	0.3	0.112%
1.0	1.5	2.9	0.45	0.05	35	0.9	0.8	0%
2.0	2.0	2.9	0.3	0.15	12	1.1	0.8	0.092%
0.4	0.4	3.0	0.85	0.15	5	1	0.2	0.749%
0.4	2.4	3.5	0.04	0.03	25	1.5	1	0.027%
2.5	3.0	3.5	0.75	0.6	3	2	1	0%

The reported results show that the gap is less than 1.5% in all cases, and practically nonexistent in many of them. This can be attributed to the fact that the optimal policy is only affected in a small number of stages, and that the travel time seems to be less sensitive to velocity decision in the affected areas, i.e., regions of slowing down.

From a practical standpoint, these findings justify the discrete-velocity scenario being our benchmark. A decision-support tool that is based on this simple model could provide recommendations in real time, something that may be challenging with a continuous-velocity model, whose computational complexity is much higher.

A second observation is that velocity oscillations persist in the continuous-velocity scenario, albeit somewhat damped, wherever they appear in the discrete-velocity experiments. Hence, on the descriptive side, this suggests that oscillations are not an artifact of our modeling choices – to consider just two velocities in our benchmark model – but rather an intrinsic characteristic of the problem.

Endogenous merging points

Finally, we analyze an extension to our benchmark model where merging points are themselves endogenous decisions of the driver, along with the (continuous) velocity. Our aim is to illustrate that velocity oscillations may persist even in this setting.

We continue with one of our main modeling constructs, the merging probability function $q(v_k^e, V_k^F, d_k)$. However, we fix the driver's belief of the future free-lane velocity V_k^F at its expected value $\bar{v}^F$. Hence, we can simplify the merging probability to $\tilde{q}(v_k^e, d_k)$. We assume $\tilde{q}(v, 0) = 0$, for all v, and that $\tilde{q}(v, d)$ is increasing in d and decreasing in v. Moreover, we assume that there exists some $\epsilon > 0$, such that

$$\tilde{q}(v, d) = 0, \quad \forall d < \epsilon.$$

This is a standard assumption in dynamic programming, e.g., see Section 11.1 of Puterman (1994), to ensure that we avoid trivial situations with infinite merging attempts. Moreover, it is also realistic: if the driver is unable to merge at some point, the chances of merging in that vicinity are practically zero.

Furthermore, our standing assumption throughout the paper is that acceleration/deceleration between merging points remains constant.

We formulate the Dynamic Program similarly to our benchmark model, whose solution provides (an optimal) set of merging points and velocities, (x_k^*, v_k^*), $k = 0, \dots, N$, with x_{N+1}^* being the blocking point. Let $d_k^* = x_{k+1}^* - x_k^*$.

We seek conditions such that $v_0^* < v_1^* > v_2^*$, which would correspond to velocity oscillations, at optimality. Fix all optimal merge-points and velocities, apart from v_1^*. We have that

$$v_1^* = \arg \min_{v_L \leq v \leq v_H} T_1(v),$$

where v_H and v_L are the upper and lower limits on feasible velocity, and

$$T_1(v) \equiv \left\{ \frac{d_0^*}{v_0^* + v} + \tilde{q}(v, d_0^*) T_1^F + \left(1 - \tilde{q}(v, d_0^*)\right) \left(\frac{d_1^*}{v_2^* + v} + \mathbb{E}\left[T^B\left(x_2^*, v_2^*\right)\right] \right) \right\}.$$

We use the Envelope Theorem in order to compute the derivative of $T_1(v)$, by viewing $x_k^*(v)$ and $v_k^*(v)$, for all $k \neq 1$, as optimized variables for every fixed v:

$$-\frac{d_0^*}{\left(v_0^* + v\right)^2} - \frac{\left(1 - \tilde{q}(v, d_0^*)\right) d_1^*}{\left(v_2^* + v\right)^2} + \tilde{q}'(v, d_0^*) \left(T_1^F - \frac{d_1^*}{v_2^* + v} - \mathbb{E}\left[T^B\left(x_2^*, v_2^*\right)\right] \right).$$

Let us assume that the vehicle is already in the Final Merging Zone, i.e., $\left(T_1^F - \frac{d_1^*}{v_2^* + v} - \mathbb{E}\left[T^B\left(x_2^*, v_2^*\right)\right]\right)$ is negative and, for simplicity, that $v_0^* = v_2^* = v_L$. We explore conditions such that the derivative of $T_1(v)$ is negative, for all $v < v_H$, which would imply that the optimal solution to the above constrained optimization problem is v_H. In other words, we require:

$$\tilde{q}'(d_0^*, v)\left(T_1^F - \frac{d_1^*}{v_2^* + v} - \mathbb{E}\left[T^B\left(x_2^*, v_2^*\right)\right]\right) < \frac{d_0^*}{\left(v_0^* + v\right)^2} + \frac{\left(1 - \tilde{q}\left(v, d_0^*\right)\right)d_1^*}{\left(v_2^* + v\right)^2}.$$

For concreteness, let us assume that $\tilde{q}(v, \cdot)$ is a linear function zero slope in the interval $[0, \epsilon]$, and slope $-\alpha$ from then on, for some $\alpha > 0$. Since $\tilde{q}'(v, \cdot)$ and $\left(T_1^F - \frac{d_1^*}{v_2^* + v} - \mathbb{E}\left[T^B\left(x_2^*, v_2^*\right)\right]\right)$ are both negative, we have that

$$\alpha\left(\frac{d_1^*}{v_2^* + v} + \mathbb{E}\left[T^B\left(x_2^*, v_2^*\right)\right] - T_1^F\right) < \frac{d_0^*}{\left(v_0^* + v\right)^2} + \frac{\left(1 - \tilde{q}\left(v, d_0^*\right)\right)d_1^*}{\left(v_2^* + v\right)^2}.$$

The right-hand side is bounded from below by $\epsilon / \left(v_0^* + v_H\right)^2$, and we obtain a sharper sufficient condition:

$$\alpha\left(\frac{d_1^*}{2v_L} + \mathbb{E}\left[T^B\left(x_2^*, v_2^*\right)\right] - T_1^F\right) < \frac{\epsilon}{\left(v_L + v_H\right)^2}.$$

Summarizing, if T_1^F and $\mathbb{E}\left[T^B\left(x_2^*, v_2^*\right)\right]$ satisfy the inequality above, then we have velocity oscillations similarly to our benchmark model, despite the fact that velocity is continuous and the merging attempt points are optimally decided.

References

Acemoglu, D., Johari, R., Ozdaglar, A., 2007. Partially optimal routing. IEEE J. Sel. Areas Commun. 25 (6), 1–13.

Acemoglu, D., Makhdoumi, A., Malekian, A., Ozdaglar, A., 2018. Informational braess' paradox: the effect of information on traffic congestion. Oper. Res. 66 (4), 893–917.

Agarwal, S., Kachroo, P., Contreras, S., Sastry, S., 2015. Feedback-coordinated ramp control of consecutive on-ramps using distributed modeling and godunov-based satisfiable allocation. IEEE Trans. Intell. Transp. Syst. 16 (5), 2384–2392.

Ahmed, K.I., 1999. Modeling Drivers' Acceleration and Lane Changing Behavior (Ph.D. thesis). MIT, Cambridge, MA.

Ahn, S., Cassidy, M.J., 2007. Freeway traffic oscillations and vehicle lane-change maneuvers. In: Transportation and Traffic Theory 2007.

Alessandri, A., Di Febbraro, A., Ferrara, A., Punta, E., 1998. Optimal control of freeways via speed signaling and ramp metering. Control Eng. Pract. 6, 771–780.

Arnott, R., de Palma, A., Lindsey, R., 1990. Economics of a bottleneck. J. Urban Econ. 27, 111–130.

Arnott, R., de Palma, A., Lindsey, R., 1991. Does providing information to drivers reduce traffic congestion? Transp. Res. A 25A (5), 309–318.

Athans, M., 1969. A unified approach to the vehicle-merging problem. Transp. Res. 3 (1), 123–133.

Awal, T., Kulik, L., Ramamohanrao, K., 2013. Optimal traffic merging strategy for communication and sensor-enabled vehicles. In: Proceedings of the 16th International IEEE Annual Conference on Intelligent Transportation Systems, pp. 1468–1474.

Barlovic, R., Esser, J., Froese, K., Knospe, W., Neubert, L., Schreckenberg, M., Wahle, J., 1999. Online traffic simulation with cellular automata. In: Brilon, Werner, Huber, Felix, Schreckenberg, Michael, Wallentowitz, Henning (Eds.), Traffic and Mobility. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 117–134.

Barlovic, Robert, Santen, Ludger, Schadschneider, Andreas, Schreckenberg, Michael, 1998. Metastable states in cellular automata for traffic flow. Eur. Phys. J. B 5 (3), 793–800.

Baykal-Gürsoy, M., Xiao, W., Ozbay, K., 2009. Modeling traffic flow interrupted by incidents. European J. Oper. Res. 195 (1), 127–138.

van Beinum, Aries, Farah, Haneen, Wegman, Fred, Hoogendoorn, Serge, 2018. Driving behaviour at motorway ramps and weaving segments based on empirical trajectory data. Transp. Res. C 92, 426–441.

Buisson, Christine, Keyvan-Ekbatani, Mehdi, Wagner, Peter, 2018. Impede autonomous vehicles merging at on-ramps? In: Conference on Intelligent Transportation Systems, 97th Annual Meeting Transportation Research Board. Washington D.C., p. 21.

Cao, W., Mukai, M., Kawabe, T., Nishira, H., Fujiki, N., 2014. Gap selection and path generation during merging maneuver of automobile using real-time optimization. SICE J. Control, Meas. Syst. Integr. 7 (4), 227–236.

Cao, W., Mukai, M., Kawabe, T., Nishira, H., Fujiki, N., 2015. Cooperative vehicle path generation during merging using model predictive control with real-time optimization. Control Eng. Pract. 34, 98–105.

Chang, G.-L., Kao, Y.-M., 1991. A realistic two-lane cellular automata traffic model considering aggressive lane-changing behavior of fast vehicle. Transp. Res. A 25A (6), 375–389.

Chen, Danjue, Ahn, Soyoung, Chitturi, Madhav, Noyce, David A, 2017. Towards vehicle automation: Roadway capacity formulation for traffic mixed with regular and automated vehicles. Transp. Res. B 100, 196–221.

Choudhury, C.F., Ben-Akiva, M.E., 2013. Modelling driving decisions: a latent plan approach. Transp. A 9 (6), 546–566.

Choudhury, Charisma F., Ben-Akiva, Moshe, Abou-Zeid, Maya, 2010. Dynamic latent plan models. J. Choice Model. 3 (2), 50–70.

Choudhury, Charisma, Ben-Akiva, Moshe, Toledo, Tomer, Rao, Anita, Lee, Gunwoo, 2006. NGSIM-Cooperative Lane Changing and Forced Merging Model. Tech. Rep. FHWA-HOP-07-096, DOT, Federal Highway Administration, Washington, DC.

Choudhury, Charisma F., Ramanujam, Varun, Ben-Akiva, Moshe E., 2009. Modeling acceleration decisions for freeway merges. Transp. Res. Rec. 2124 (1), 45–57.

Chowdhury, D., Santen, L., Schadschneider, A., 2000. Statistical physics of vehicular traffic and some related systems. Phys. Rep. 329 (4), 199–329.

Como, G., Lovisari, E., Savla, K., 2016. Convexity and robustness of dynamical traffic assignment for control of freeway networks. Transp. Res. Part B 91, 446–465.

Correa, J., Schultz, A., Stier-Moses, N., 2004. Selfish routing in capacitated networks. Math. Oper. Res. 29 (4), 961–976.

Daganzo, C.F., 2002a. A behavioral theory of multi-lane traffic flow part I: Long homogeneous freeway sections. Transp. Res. B 36, 131–158.

Daganzo, C.F., 2002b. A behavioral theory of multi-lane traffic flow part II: Merges and the onset of congestion. Transp. Res. B 36, 159–169.

Del Castillo, J.M., Benítez, F.G., 1995. On the functional form of the speed-density relationship: General theory. Transp. Res. B 29 (5), 373–389.

Duret, A., Bouffier, J., Buisson, C., 2010. Onset of congestion from low-speed merging maneuvers within free-flow traffic stream. Transp. Res. Record: J. Transp. Res. Board 2188, 96–107.

Ebersbach, A., Schneider, J.J., 2004. Two-lane traffic with places of obstruction to traffic. Int. J. Modern Pys. C 15 (4), 535–544.

Erdmann, Jakob, 2014. Lane-changing model in SUMO. In: SUMO2014, Reports of the DLR-Institute of Transportation SystemsProceedings, Vol. 24. Deutsches Zentrum für Luft- und Raumfahrt e.V., pp. 77–88, URL https://elib.dlr.de/89233/,

Evans, D.H., Herman, R., Weiss, G.H., 1964. The highway merging and queuing problem. Oper. Res. 12 (6), 832–857.

Flotterod, G., Bierlaire, M., Nagel, K., 2011. Bayesian demand calibration for dynamic traffic simulations. Transp. Sci. 45 (4), 541–561.

Fritzsche, H.-T., 1994. A model for traffic simulation. Transp. Eng. Contrib. 5 (4), 317–321.

Gipps, P.G., 1986. A model for the structure of lane-changing decisions. Transp. Res. B 35, 107–120.

Greenshields, B.D., 1934. A study in highway capacity. Highway Res. Board Proc. 14, 448–477.

Gregoire, J., Qian, X., Frazzoli, E., de la Fortelle, A., Wongpiromsarn, T., 2015. Capacity-aware backpressure signal control. IEEE Trans. Control Netw. Syst. 2 (2), 164–173.

Grillo, L., Datta, T., Hartner, C., 2008. Dynamic late lane merge system at freeway construction work zones. Transp. Res. Record: J. Transp. Res. Board 2055, 3–10.

Han, Y.-S., Ko, S.-K., 2012. Analysis of a cellular automaton model for car traffic with a junction. Theoret. Comput. Sci. 450, 54–67.

Hawkes, A.G., 1968. Gap acceptance in road traffic. J. Appl. Probab. 5, 84–92.

Heidemann, D., 2001. A queueing theory model of nonstationary traffic flow. Transp. Sci. 35 (4), 405–412.

Helbing, D., 2001. Traffic and related self-driven many-particle systems. Rev. Modern Phys. 73, 1067–1141.

Hidas, P., 2002. Modelling lane changing and merging in microscopic traffic simulation. Transp. Res. C 10, 351–371.

Hidas, P., 2005. Modelling vehicle interactions in microscopic simulation of merging and weaving. Transp. Res. 13 (1), 37–62.

Hounsell, N.B., Barnard, S.R., McDonald, M., Owens, D., 1992. An Investigation of Flow Breakdown and Merge Capacity on Motorways. TRL Contractor Report 338, Transport and Road Research Laboratory, Crowthorne, Berkshire, UK.

Iordanidou, G.-R., Roncoli, C., Papamichail, I., Papageorgiou, M., 2015. Feedback-based mainstream traffic flow control for multiple bottlenecks on motorways. IEEE Trans. Intell. Transp. Syst. 16 (2), 610–621.

Jain, R., Smith, J.M., 1997. Modeling vehicular traffic flow using M/G/C/C state dependent queueing models. Transp. Sci. 31 (4), 324–336.

Jin, W.-L., 2010. A kinematic wave theory of lane-changing traffic flow. Transp. Res. B 44, 1001–1021.

Kachroo, P., Li, Z., 1997. Vehicle merging control design for an automated highway system. In: Proceedings of the IEEE Conference on Intelligent Transportation Systems, pp. 224–229.

Kesting, A., Treiber, M., Helbing, D., 2007. General lane-changing model MOBIL for car-following models. Transp. Res. Record: J. Transp. Res. Board 1999, 86–94.

Keyvan-Ekbatani, Mehdi, Knoop, Victor L., Daamen, Winnie, 2016. Categorization of the lane change decision process on freeways. Transp. Res. C 69, 515–526.

Kita, H., 1999. A merging-giveway interaction model of cars in a merging section: a game theoretic analysis. Transp. Res. Part A 33 (3–4), 305–312.

Knoop, V.L., Hoogendoorn, S.P., Shiomi, Y., Buisson, C., 2012. Quantifying the number of lane changes in traffic: Empirical analysis. Transp. Res. Rec. 2278 (1), 31–41.

Kondyli, Alexandra, Elefteriadou, Lily, 2011. Modeling driver behavior at freeway–ramp merges. Transp. Res. Rec. 2249 (1), 29–37.

Kotsialos, A., Papageorgiou, M., 2004. Nonlinear optimal control applied to coordinated ramp metering. IEEE Trans. Control Syst. Technol. 12 (6), 920–933.

Laval, J.A., Daganzo, C.F., 2006. Lane-changing in traffic streams. Transp. Res. B 40 (3), 251–264.

Le, T., Kovacs, P., Walton, N., Vu, H., Andrew, L., Hoogendorn, S., 2015. Decentralized signal control for urban road networks. Transp. Res. C 58, 431–450.

Lee, G., 2006. Modeling Gap Acceptance at Freeway Merges (Ph.D. thesis). Massachusetts Institute of Technology, Dept. of Civil and Environmental Engineering.

Liu, R., Hyman, G., 2012. Modelling motorway merge: the current practice in the UK and towards establishing general principles. Transp. Policy 24, 199–210.

Lopez, Pablo Alvarez, Behrisch, Michael, Bieker-Walz, Laura, Erdmann, Jakob, Flötteröd, Yun-Pang, Hilbrich, Robert, Lücken, Leonhard, Rummel, Johannes, Wagner, Peter, ner, Evamarie Wieß, 2018. Microscopic traffic simulation using SUMO. In: The 21st IEEE International Conference on Intelligent Transportation Systems. IEEE, URL https://elib.dlr.de/124092/.

Mauch, M., Cassidy, M.J., 2002. Freeway Traffic Oscillations: Observations And Predictions. Emerald Group Publishing Limited, pp. 653–673, chap. 32.

McCoy, P.T., Pesti, G., 2001. Dynamic late merge-control concept for work zones on rural interstate highways. Transp. Res. Record 1745, 20–26.

McNeil, D.R., Smith, J.T., 1969. A comparison of motorist delays for different merging strategies. Transp. Sci. 3 (3), 239–254.

McQueen, J., Miller Jr., R.G., 1960. Optimal persistence policies. Oper. Res. 8 (3), 362–380.

Miller, A.J., 1961. A queueing model for road traffic flow. J. R. Stat. Soc. Ser. B Stat. Methodol. 23 (1), 64–90.

Nagel, K., Schreckenberg, M., 1992. A cellular automaton model for freeway traffic. J. Physique I 2 (12), 2221–2229.

Ngoduy, D., Lee, S., Treiber, M., Keyvan-Ekbatani, M., Vu, H.L., 2019. Langevin method for a continuous stochastic car-following model and its stability conditions. Transp. Res. C 105, 599–610.

Nilsson, J., Silvlin, J., Brannstrom, M., Coelingh, E., Fredriksson, J., 2016. If, when, and how to perform lane change maneuvers on highways. IEEE Intell. Transp. Syst. Mag. 8 (4), 68–78.

de Palma, A., Ben-Akiva, M., Lefevre, C., Litinas, N., 1983. Stochastic equilibrium model of peak period traffic congestion. Transp. Sci. 17, 430–453.

Pasquale, C., Papamichail, I., Roncoli, C., Sacone, S., Siri, S., Papageorgiou, M., 2015. Two-class freeway traffic regulation to reduce congestion and emissions via nonlinear optimal control. Transp. Res. C 55, 85–99.

Pipes, L.A., 1967. Car-following models and the fundamental diagram of road traffic. Transp. Res. 1 (1), 21–29.

Puterman, M.L., 1994. Markov Decision Processes: Discrete Stochastic Dynamic Programming, 1st edn. John Wiley & Sons, Inc., USA.

Raravi, G., Shingde, V., Ramamritham, K., Bharadia, J., 2007. Merge algorithms for intelligent vehicles. In: Next Generation Design and Verification Methodologies for Distributed Embedded Control Systems. pp. 51–65.

Rickert, M., Nagel, K., Schreckenberg, M., Latour, A., 1996. Two lane traffic simulations using cellular automata. Physica A 231 (4), 534–550.

Rios-Torres, J., Malikopoulos, A.A., 2017. A survey on the coordination of connected and automated vehicles at intersections and merging at highway on-ramps. IEEE Trans. Intell. Transp. Syst. 18 (5), 1066–1077.

Rios-Torres, J., Malikopoulos, A.A., Pisu, P., 2015. Online optimal control of connected vehicles for efficient traffic flow at merging roads. In: Proceedings of the 18th International IEEE Annual Conference on Intelligent Transportation Systems, pp. 2432–2437.

Schakel, Wouter J., Knoop, Victor L., van Arem, Bart, 2012. Integrated lane change model with relaxation and synchronization. Transp. Res. Rec. 2316 (1), 47–57.

Shalev-Shwartz, S., Shammah, S., Shashua, A., 2016. Safe, multi-agent reinforcement learning for autonomous driving. ArXiv, 1610.03295, URL http://arxiv.org/abs/1610.03295.

Shalev-Shwartz, S., Shammah, S., Shashua, A., 2017. On a formal model of safe and scalable self-driving cars. CoRRabs/1708.06374, URL http://arxiv.org/abs/1708.06374.

Tarko, A.P., Shamo, D., Wasson, J., 1999. Indiana lane merge system for work zones on rural freeways. J. Transp. Eng. 125 (5), 415–420.

Tobita, K., Naito, Y., Nagatani, T., 2012. Traffic flow merging and bifurcating at junction on two-lane highway. World J. Mech. 2 (04), 203–251.

Toledo, T., Koutsopoulos, H.N., Ben-Akiva, M.E., 2003. Modeling integrated lane-changing behavior. Transp. Res. Rec. 1857 (1), 30–38.

Treiber, Martin, Hennecke, Ansgar, Helbing, Dirk, 2000. Congested traffic states in empirical observations and microscopic simulations. Phys. Rev. E 62 (2), 1805.

Treiber, Martin, Kesting, Arne, 2018. The intelligent driver model with stochasticity–new insights into traffic flow oscillations. Transp. Res. B 117.

Underwood, R.T., 1961. Speed, Volume and Density Relationships. Yale Bur. Highway Traffic, New Haven, Connecticut.

Vickrey, W.S., 1969. Congestion theory and transport investment. Amer. Econ. Rev. 59, 251–261.

Wagner, P., Nagel, K., Dietrich, E.W., 1997. Realistic multi-lane traffic rules for cellular automata. Physica A 234 (3–4), 687–698.

Wan, Xia, Jin, Peter J, Gu, Haiyan, Chen, Xiaoxuan, Ran, Bin, 2017. Modeling freeway merging in a weaving section as a sequential decision-making process. J. Transp. Eng. Part A 143 (5), 05017002.

Wardrop, J.G., 1952. Some theoretical aspects of road traffic research. Proc. Inst. Civ. Eng. 1 (3), 325–362.

Weiss, G.H., Maradudin, A.A., 1962. Some problems in traffic delay. Oper. Res. 10, 74–104.

Yang, Q., Koutsopoulos, H.N., 1996. A microscopic traffic simulator for evaluation of dynamic traffic management systems. Transp. Res. C 4 (3), 113–129.

Zhang, J., Li, X., Wang, R., Sun, X., Cui, X., 2012. Traffic bottleneck characteristics caused by the reduction of lanes in an optimal velocity model. Physica A 391 (7), 2381–2389.

Zheng, Z., 2014. Recent developments and research needs in modeling lane changing. Transp. Res. B 60, 16–32.

Zheng, Z., Ahn, S., Chen, D., Laval, J., 2013. The effects of lane-changing on the immediate follower: Anticipation, relaxation, and change in driver characteristics. Transp. Res. C 26, 367–379.

Zheng, Z., Ahn, S., Monsere, C.M., 2010. Impact of traffic oscillations on freeway crash occurrences. Accid. Anal. Prev. 42 (2), 626–636.

Zheng, Fangfang, Liu, Can, Liu, Xiaobo, Jabari, Saif Eddin, Lu, Liang, 2020. Analyzing the impact of automated vehicles on uncertainty and stability of the mixed traffic flow. Transp. Res. C 112, 203–219.