Managing algoritm-assisted drivers

Dmitrii Tikhonenko*

JOB MARKET PAPER

Click here for the latest version

October 10, 2019

Abstract

Why do traffic jams and stop-and-go waves happen? Is there a way to prevent them? Traffic theory models ascribe the fundamental cause of such phenomena to exogenous fluctuations originating from human driving behavior. But what if vehicles are controlled or assisted by algorithms? Would these just go away? I study a simple, yet very common forced-merging scenario: a two-lane road segment that has one lane blocked unexpectedly, say, due to an accident or construction. All cars on the blocked lane need to merge to the free lane in order to continue their routes. Motivated by the recent advent of algorithm-assisted driving, I assume that drivers are rational, self-interested agents, wishing to minimize their individual travel times, deciding (a) at what velocity to move; and (b) whether to merge to the free lane, given the opportunity (gap). Moving at higher velocities on the blocked lane reduces the travel time but also the chance of finding a large enough gap to merge, so blocked-lane drivers are trading off travel time vs. risk of not being able to merge. I analyze a dynamic programming formulation of the problem with a single merging driver, and characterize the optimal policy, which turns out to be a multi-threshold one with a surprising structure: in the presence of uncertainty regarding merging to the target lane, it may be optimal for a driver–in certain regions–to oscillate between high and low velocities while attempting to merge. Hence, the origin of traffic oscillations need not be purely "behavioral" but can arise endogenously as the outcome of optimization. I show how a central planner can set velocity limits for avoiding such oscillations, and test our policies via extensive discrete event simulations with multiple merging vehicles.

Keywords: traffic bottleneck; optimal merging; rational drivers; dynamic programming.

1 Introduction

Optimizing traffic flow can reduce environmental pollution, improve quality-of-life for commuters, and save billions of dollars in lost productivity. Traffic and its control is therefore of significant interest and intensely studied in the areas of Traffic Engineering, Economics, and

*Department of Economics, Universitat Pompeu Fabra. Email: dmitrii.tikhonenko@upf.edu. I am grateful to my advisors Kalyan Talluri and Mihalis Markakis for their constant guidance and support. I thank participants of 1st EUROYoung Workshop in Seville and 3rd ICITT Conference in Amsterdam for their valuable comments and discussion.

Operations Research. The nature of road traffic, however, poses significant modeling and technical challenges: it is highly dynamic and uncertain, with a large number of autonomous and self-interested agents of varying skills and attitudes, with the decisions of one agent potentially affecting the actions of many others. Moreover, it is subject only to broad fixed traffic guidelines which by necessity have to be applicable for vastly changing driving conditions on the same stretch of road.

Despite many decades of research, the fundamental reasons for certain traffic phenomena, such as stop-and-go waves and oscillations are still not completely understood; see Zheng (2014) for a recent critique. Indeed, models of traffic that aim to derive analytical insights tend to be macroscopic in nature, geared towards capturing the aggregate phenomena of the system at hand with an exogenous modeling element (driver behavior) that causes fluctuations. Examples include car-following, fluid and gas-kinetic models proposed by physicists, queueing models from Operations Research, and equilibrium models in Economics.

In contrast, the literature capturing the movements and interactions of every car at a microscopic level, e.g., particle-hopping and cellular automata models, typically relies on simulations with stochastic elements. The behavior of drivers in such simulations is governed by certain sets of rules, usually calibrated vis-a-vis real traffic and observations of driver behavior. The goal of the simulations is to evaluate the performance of some policy measured by macro-level traffic metrics but suffer from the ad-hoc nature of the rules and specificity of most simulation studies. (A more detailed review of the existing literature follows in §2). Modeling the independent decision-making process and the incentives of autonomous agents and deriving any sort of conclusive insight is recognized as exceptionally challenging.

That said, there is a compelling new argument in favor of modeling and studying traffic control at the microscopic policy level: the advent of sensors, vehicle-to-vehicle communication, and driverless car technologies. New communication technologies such as V2V (vehicle to vehicle) and V2I (vehicle to infrastructure) enable the exchange of information, control and decisions in real-time. Both regulators and car manufacturers seem ready to adapt these new standards in the coming years (see New York Times (2016), DOT (2018)). It even seems plausible that a central planner can coordinate traffic (akin to automated air-traffic control), or at least provide real-time information, with the new ability to communicate to vehicles. Algorithms used in self-driving cars are more likely than humans to follow rules and have clearly defined objectives, so they could potentially reach a mutually beneficial consensus with other vehicles through V2V communication in real-time. All these developments raise the need for a better understanding, via models and theoretical analysis, of optimal policies that take drivers' objectives and self-interested nature into consideration.

On a separate but related note, there are evidences that improvements at the macro-level, such as opening new lanes or highways, repeatedly fails to reduce congestion; Duranton and Turner (2011). Driverless cars, coupled with well designed policies and efficient algorithms at the car level, could potentially do much better; New York Times (2017).

Upon some reflection however, it should become apparent that in a new transportation paradigm where driving is assisted by algorithms, either fully or partially, the real challenge is the management of algorithmic-assisted vehicular traffic. Unlike humans who often rely on visual, cultural, or behavioral cues, algorithms have more concrete objective functions and need protocols to resolve the navigation and safety issues on the road. A new direction of research is emerging to tackle this challenge that combines ideas from Computer Science, Operations Research, Microeconomics, Public Policy and Management; for instance, see Shalev-Shwartz et al. (2016), Shalev-Shwartz et al. (2017).

This paper studies traffic control from a *management* point of view: what should the protocols be for overtaking, passing, and merging? What policies lead to fair, socially optimal, or efficient outcomes? Why should the cars follow a prescribed policy, if it is not in their interest to do so? While the technology behind driverless cars is advancing at a breath-taking pace, much is still unknown about such issues of optimization, incentives, and collective decision-making by heterogenous agents, and the resulting outcomes.

Specifically, are fluctuations that are known to lead to inexplicable traffic jams and stop-and-go phenomena due solely to random driver behavior, as modeled in most macroscopic traffic models? Would they go away if vehicles are managed by algorithms?

Motivated by a need to understand such questions, in this paper I analyze a simple stylized model of a very common traffic situation: due to an accident or construction work, one of the two lanes of a road segment is blocked unexpectedly. All drivers that happen to be on the blocked lane need to merge to the free lane in order to continue with their routes—called a forced-merging scenario. Assuming that drivers are rational[1] and self-interested, our goal is to find a lane-merging and velocity-control policy that minimizes their expected travel time.

Drivers on the blocked lane know they have to merge, and each driver has the following decisions to make: (i) the velocity to move; (ii) whether to merge to the free lane, given the opportunity. Information is not perfect and they have only a stochastic view of the state of the free lane. What makes the problem interesting is that moving at a high velocity on the blocked lane reduces immediate travel time, but also reduces the chance of finding a large enough gap to merge increasing the risk of getting stuck. So the drivers are trading off instantaneous travel time against probability of merging. I give a Dynamic Programming (DP) formulation of the problem and characterize the optimal policy, which turns out to be a multi-threshold one with a non-trivial and surprising structure: in the presence of uncertainty regarding the future state of the target lane at the time of merging, it may be optimal for a driver—under certain parameter regimes—to oscillate between high and low velocities while attempting to merge, i.e., rational, optimizing, driving behavior may lead to oscillatory traffic patterns; in other words, fluctuations can arise endogenously, rather than due to random driver behavior.

Previous theoretical models include an exogenous random variable to incorporate such fluctuations and explain traffic jams and stop-and-go waves. Oscillations in velocities of a driver are known to cause traffic jams and accidents. Empirical and observational research seems to indicate that lane-changing behavior may be one of the causes of such oscillations. Hence, on the *descriptive* side, I provide an alternative, optimization-based explanation for traffic oscillations, a phenomenon that has perplexed both laymen and researchers alike. On the *normative* side, our paper sets the foundations for decision-support systems for algorithm-assisted drivers; I give conditions on speed limits so that such regions of oscillations are eliminated, improving traffic flow in our simulations.

The remainder of the paper is organized as follows. In §2 I provide a detailed review of the related literature. §3 is devoted to the analytical treatment of the problem: §3.1 presents a stylized model of the merging process, together with a discussion of our modeling assumptions and a DP formulation of the problem; §3.2 contains our analytical results, which, among others, showcase how rational, self-interested merging can lead to "irrational" traffic patterns. §4 adopts the viewpoint of a traffic manager: how to choose optimally the speed limits for the blocked lane, so as to minimize rational drivers' travel times and eliminate velocity oscillations. Finally, §5 provides an evaluation of the proposed approach, as well as comparisons to several intuitive heuristics, through extensive discrete-event simulations. In §6 I conclude the paper with a brief discussion

[1] I use the word *rational* less in the formal economic sense than to distinguish from typical driving behavior which is very varied and occasionally inexplicable.

of our findings and directions for future work.

2 Literature Review

Merging rules and discretionary lane-changing have received considerable attention in the research communities that study the problems of road traffic: Operations Research, Physics, Economics, Systems, Control, and Traffic Engineering.

Before I present the relevant literature in detail, let us attempt a brief synopsis of the various approaches, so as to position the present paper in broad terms. Generally speaking, traffic modeling follows either a macroscopic approach, studying traffic at the population/fluid level, or a microscopic approach, studying the behavior of individual drivers[2]. The former provides valuable insights into the evolution and equilibrium behavior of the large-scale, complex systems at hand, but fails to capture the incentives and behavior of individual drivers; moreover, many of the differential equations introduce an exogenous and independent random variable ("driver behavior") to explain fluctuations. The latter approach studies more realistic models of traffic, but their complexity prevents a theoretical analysis, so performance evaluation is typically done through micro-simulations, usually in the cellular automata framework.

Traffic research flourished in the early days of Operations Research, often viewed through the lens of queueing theory, as in Miller (1961). The merging problem in particular was the topic of early influential papers, 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. However, traffic is modeled as an exogenous stochastic process, and the individual behavior and incentives of drivers is absent from these early works. After an inexplicably long hiatus, the community is picking up this important topic, as in Jain and Smith (1997), Heidemann (2001), based on queueing theory, and Gregoire et al. (2015), Le et al. (2015) and Como et al. (2016), from a control and optimization perspective.

Merging is essentially an optimal stopping problem with similarities to the classical "parking problem" [3]: a driver moves along a street looking for space to park her car, as close to her destination as possible. Availability of slots is uncertain and known only probabilistically. The optimal parking policy is a threshold one: the driver parks in the first available spot, but only after passing over slots for a certain distance; see McQueen and Miller Jr. (1960). The main difference of our setting 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 space to merge. While the optimal policy in our case is still of threshold type, a multi-threshold one, our analysis reveals that its structure is surprisingly richer than that of the classical parking problem, allowing for both smooth "monotonic" behavior and velocity oscillations, depending on the parameters of the problem.

The literature in Economics emphasizes the behavior and incentives of rational self-interested drivers, focusing on the resulting equilibria. Notable are the studies on the inefficiencies of traffic networks with selfish agents, e.g., the Braess paradox; see Wardrop (1952), Vickrey (1969),

[2]A smaller body of work is classified as *mesoscopic* that attempts to describe microscopic dynamics as a function of macroscopic parameters.

[3]A somewhat lesser known problem in control theory is the "rocket-rail car problem," of parking a rail car with rocket engines on both ends. The goal is to fire the rockets to make the car stop at a precise point in the least amount of time. While this problem is closer to the one I study in this work, in the sense that it has a velocity element to it, I were unable to translate the results to our setting.

de Palma et al. (1983), Arnott et al. (1990), Arnott et al. (1991), Correa et al. (2004), Lago and Daganzo (2007), Acemoglu et al. (2007) and Acemoglu et al. (2018). All these papers can be said to be 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.

The number of distinct traffic models and theories proposed is too large for us to even give an overview here. I refer the reader to a relatively recent survey due to Helbing (2001) from a physicist point of view and Chowdhury et al. (2000) from the cellular automata perspective. I recommend also a more recent survey of Zheng (2014) that nicely summarizes existing contributions, and also new modeling approaches proposed in Shalev-Shwartz et al. (2016) and Shalev-Shwartz et al. (2017) with algorithmic driving in mind.

I highlight next the literature focusing specifically on the merging problem.

2.1 Literature specific to Lane Changing

Over the last two decades, lane-changing has gained prominence in the Traffic Science 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.

The Systems and Control community has viewed traffic, specifically merging at highway on-ramps, as an optimal control problem, where every car is controlled by a central planner. 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 and Malikopoulos (2017a). While providing a useful benchmark, this viewpoint, again, fails to capture the incentives of individual drivers. A decentralized approach, where cars are not centrally controlled but 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). I note that there is also a large number of less principled, heuristic approaches to merging in this strand of literature; I refer the interested reader to the excellent survey paper Rios-Torres and Malikopoulos (2017b).

Many empirical works have been conducted on how drivers behave at merging points, e.g., Chang and Kao (1991), Hounsell et al. (1992), Hidas (2005), and Liu and Hyman (2012), Knoop et al. (2012); as well as modeling works on merging and lane-changing proposed, e.g., Gipps (1986), Kita (1999), Daganzo (2002a), Daganzo (2002b), Hidas (2002), Jin (2010), Laval and Daganzo (2006) and Zhang et al. (2012). Choudhury et al. (2007), Choudhury and Ben-Akiva (2013) and Kesting et al. (2007) propose sophisticated merging decision structures and different merging regimes, including potential cooperation with cars in the target lane. A notable difference with our model, though, is the lack of a velocity decision. I refer the reader to Sun and Elefteriadou (2010) (which does have a velocity element and validated by real data), Rahman et al. (2013), and Zheng (2014) for a broader review of lane-change models.

Of particular interest and relevance to this paper are some real-world observational studies that link lane-changing behavior to oscillations between high and low velocities. Zheng (2014), Ahn and Cassidy (2007), and Mauch and Cassidy (2002) conduct such studies and point out that their data does not support the classical car-following explanations, but instead oscillations are strongly linked to lane-changing behaviour. However, the oscillations they study are on the free lane and a result of the disturbances that merges create. In contrast, our results are surprising in that drivers on the blocked lane may oscillate their velocities as a result of optimizing behavior. Zheng et al. (2010) reports a strong correlation between oscillations and the occurrence of accidents.

Sheu (2013) and Wei et al. (2000) develop simulation models backed by video evidence of lane-changing behaviour. The thesis of Ahmed (1999) proposes a logistic-regression type model to determine the probability of a driver choosing to merge. Toledo et al. (2003) proposes an integrated lane-changing model that works for mandatory and discretionary considerations based on a gap-acceptance model.

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 a theoretical analysis 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 predficting the effect of lane change using kinematic-wave theory, with a future plan to use these results for better decisions in traffic management.

Finally, a related strand of literature lies in the intersection of Traffic Engineering and Computer Science, where lane-merging problems are studied through simulations at a microscopic level. Chang and Kao (1991), Nagel and Schreckenberg (1992), Fritzsche (1994), Rickert et al. (1996), and Wagner et al. (1997), use Cellular-Automata micro-simulation models, setting driver-level rules for lane changing on multi-lane highways. Even more detailed micro-simulations have appeared in the literature, e.g., the sophisticated merging rules in MITSIM; see Yang and Koutsopoulos (1996). Micro-simulations are also used for traffic studies at the network level, e.g., Flotterod et al. (2011) considers the problem of demand estimation in traffic networks. From this literature, let us highlight Ebersbach and Schneider (2004) and Han and Ko (2012), which focus on merging into a highway with a blocked lane and mandatory on-ramp merging, respectively. Similar in spirit, although strictly speaking not using Cellular Automata, is the approach in Tobita et al. (2012), which studies merging and bifurcations in a two-lane highway. I note that, again, there is no clear control/optimization dimension, and no theoretical analysis in the aforementioned works.

3 Modeling and analysis of merging

Consider a two-lane road, with traffic moving in the same direction on both lanes. One lane is blocked unexpectedly due to an accident, construction, or maintenance work, while traffic on the other lane continues to flow freely; see Figure 1. I refer to the two lanes as B (Blocked) and F (Free), respectively. Every car that happens to be on lane B when the blockage takes place, needs to merge to lane F at some place before the blockage point, i.e., there is a scenario of *mandatory lane change*.

I start by focusing on the optimal merging and velocity control problem of a single car on lane B, say the first car, moving at only two possible velocities, a high and a low one. Our goal is to develop a tractable model that deepens our understanding of rational (vis-a-vis "behavioral") merging, provides merging and velocity recommendations to drivers, and helps traffic regulators determine the optimal speed limits that guarantee a smooth and efficient traffic flow. Later, I consider a more realistic setting, e.g., with several cars on both lanes, which I analyze through extensive discrete-event simulations. Our simulations show that the proposed policy is superior to several heuristics, in terms of both flow (throughput) and total travel time (delay).

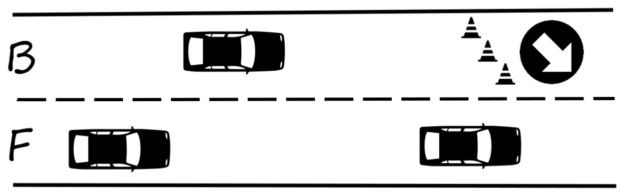

Figure 1: The illustration of a blocked lane merging scenario

3.1 A model of merging

I discretize the road segment in cells of equal length (normalized to one) that I call *stages*, starting at stage 1, until the blockage point/*bottleneck* at stage N. I use the superscripts B and F to denote stages on the blocked and free lane, respectively.

I focus on the behavior of a single car on lane B, whose goal is to merge to lane F and get through the bottleneck in the shortest possible time. The timing of decision making[4] is also discrete in our model. The driver on lane B has two decisions to make at the beginning of stage $k \in \{1, \ldots, N-1\}$: (a) whether to merge to lane F at the end of stage k, if there is a sufficient gap to do so; and (b) whether to accelerate, decelerate, or keep a constant speed during stage k, effectively determining the velocity of the car at the end of stage k.

Two important assumptions underlie our model: (a) that the driver on lane B cannot predict accurately the "state" of lane F when she reaches the end of the stage; equivalently, she cannot predict with certainty the existence of a sufficient gap for her to merge at the end of the stage; (b) that decisions and actions are not taken and executed instantaneously, instead drivers (both human as well as algorithmic) plan a bit ahead to execute an action. In our case, the driver decides to merge at the beginning 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 as well as action, is certainly realistic for humans, and will be reduced but unlikely to go away completely even in fully automated driving scenarios.

Free lane traffic is assumed to move at an average velocity $\bar{v}_F$ that is known to the driver on lane B; and these statistics do not change in the short run. I assume that once a car 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 $\bar{v}_F$ and where the merging occurs.

A merging car creates a local disturbance in the flow of the target lane, in our case lane F. This disturbance affects one or more cars on the target lane, as well as the merging car itself. The phenomenon, termed "relaxation" in the literature, is well studied and understood empirically; see Zheng et al. (2013) for a recent review. I do not attempt to model the details of the relaxation process as our focus is on the merging car and its objective of minimizing its total travel time. Hence, for simplicity and tractability, I capture the aggregate impact of relaxation on the merging car by a *merging penalty*, measured in travel time, that depends on both the velocity of the merging car as well as the velocity of the target lane.

Let v_k^s and v_k^e denote the velocity of the car at the start and at the end of stage k, respectively. I assume that these velocities can only take two values, v_H and v_L, with $v_H > v_L \geq \bar{v}_F$ (i.e., there

[4] Our discrete time, discrete-space modeling is closest to the Cellular Automata body of traffic engineering literature.

is a relatively congested free lane, typical in forced merging scenarios).

The decision on the velocity, which takes place at the beginning of each stage, affects the total travel time in three ways: a higher velocity (a) reduces the travel time on lane B, as the driver moves faster; (b) reduces the probability of merging at the end of the stage, as traffic gap-acceptance theory implies that she would need a larger gap to merge; (c) decreases the merging penalty, if she merges, as the relaxation literature suggests that the whole process becomes smoother given a larger merging gap.

I model this as follows[5]: let Y_k be a random variable that takes the value $y_k = 1$ if a sufficiently large gap is realized at the respective stage of lane F; and $y_k = 0$, otherwise. I capture the driver's uncertainty about the evolution of the state of lane F by modeling Y_k as a Bernoulli random variable, whose parameter is the end-velocity of the stage, $q(v_k^e)$, where $q(v_k^e) = q_H$ if $v_k^e = v_H$, and $q(v_k^e) = q_L$ if $v_k^e = v_L$.

I assume that $q_L > q_H > 0$, i.e., the probability of merging is higher if the velocity is lower. This is motivated by the fact that the merging probability is typically higher if the difference between the velocities in the two lanes is small, e.g., see Lee (2006); and in our case v_L is closer to $\bar{v}_F$, as our main interest is in scenarios where the free lane is relatively dense. I discuss both empirical and modeling aspects of merging, including the gap-acceptance theory, in more detail in Online Appendix B

The merging penalty depends on the velocity of the merging car at the end of stage k and on the velocity of lane F. Since the average velocity in lane F is assumed to be constant, I suppress the latter dependence and denote the merging penalty[6] by $C(v_k^e)$.

I let $C(v_k^e) = c_L$ and $C(v_k^e) = c_H$, if $v_k^e = v_L$ and $v_k^e = v_H$, respectively. I assume that $c_H < c_L$, since, on average, merging requires a larger gap if $v_k^e = v_H$, and conditional on this, the relaxation process is smoother overall. A more detailed discussion of relaxation and merging penalties can be found in the Online Appendix C.

At the beginning of stage k, the driver on lane B needs to decide whether to merge at the end of stage k if there is an opportunity to do so, i.e., if Y_k has a realization of 1. I use a binary decision variable, u_k, which takes the value 1 if the driver decides to merge, and the value 0 otherwise. If the driver reaches the bottleneck without having merged to lane F, she incurs an expected late-merging penalty of P time units, and then exits the bottleneck. This penalty models the fact that the car 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 and the solution to the driver's problem becomes trivial. In the remainder of the paper, I assume that $P \gg 1 \big/ \min\big\{v_L, \bar{v}_F\big\} + c_L$. (Recall that the length of each stage is normalized to be 1.)

I formulate the merging problem of an optimizing driver (i.e., optimal merging and velocity control) within the framework of Dynamic Programming (DP), with the following primitives.

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

Control: The control/decision at stage k, assuming the car is still on lane B, consists of two components. The first component represents the driver's intention to merge

[5]An alternate interpretation is that the flow on the F lane is Poisson with a rate $\bar{v}_F$ and the B-lane car can merge if the inter-arrival distance is above a certain threshold. However the Bernoulli model I state leads to a cleaner insight and interpretation with less jargon.

[6]$C(v_k^e)$ can be interpreted the expected delay in travel time that the merging driver experiences due to relaxation, where the expectation is with respect to her beliefs about the future state of lane F and the relaxation process dynamics.

at the end of stage k, given an opportunity, $u_k \in \{0,1\}$. The second component is the velocity to achieve at the end of stage k, $v_k^e \in \{v_L, v_H\}$. If the car has already merged to lane F at stage k, then there is no decision to be made.

Uncertainty: The uncertainty concerns the existence of a sufficient gap on lane F, at the end of stage k. I capture this through the Bernoulli random variable Y_k, whose parameter $q(\cdot)$ is a function of the velocity at beginning of stage k.

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

The vehicle reaches stages $1 \ldots N$ sequentially. The time that it takes to traverse stage k is equal to $\frac{2}{(v_k^s + v_k^e)}$, if the car is still on lane B; and $\frac{1}{\bar{v}_F}$, if it has already merged to lane F. In other words, as physics dictates (for a constant rate of acceleration), I assume that the time to traverse a stage is reciprocal to the car's average velocity during that stage.

I denote by $T_k^B(v_k^s)$ the optimal expected remaining travel time (i.e., the value function of the DP), if at the beginning of stage k the car is on lane B and moves at velocity v_k^s. Similarly, T_k^F represents the cost-to-go if the car is on lane F. I are interested in calculating $T_1^B(v_1^s)$ and, along the way, in obtaining the optimal merging and velocity control policy.

This can be done via the following DP recursion:

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

with the boundary condition

$$T_N^B(v_N^s) = P,$$

where

$$T_k^F = \frac{N-k+1}{\bar{v}_F}.$$

3.2 "Irrational" Traffic Patterns

In this section, I study the region and parameters where a seemingly odd fluctuating behavior emerges—for a simple underlying optimization reason.

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

Proposition 1. *Consider the driver's problem with velocity decisions in $\{v_L, v_H\}$. 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, I 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

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

Rational merging can be viewed as an optimal stopping problem, where the driver incurs a lump-sum cost $T^F_{k+1} + C(v^e_k)$ if she merges successfully at the end of stage k, and after which the whole process terminates prematurely. In that light, $B_k\left(v^s_k\right)$ 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 merging problem with velocity decisions in* $\{v_L, v_H\}$. *An optimal policy has the following multi-threshold structure:*

1. $(u^*_k, v^{*e}_k) = (0, v_H)$ *if and only if* $T^B_{k+1}(v_H) \leq T^F_{k+1} + c_H$;
2. $(u^*_k, v^{*e}_k) = (1, v_H)$ *if and only if* $T^B_{k+1}(v_H) \geq T^F_{k+1} + c_H$ *and* $\Delta T_{k+1} \leq B_k(v^s_k)$;
3. $(u^*_k, v^{*e}_k) = (1, v_L)$ *if and only if* $T^B_{k+1}(v_H) \geq T^F_{k+1} + c_H$ *and* $\Delta T_{k+1} \geq B_k(v^s_k)$.

So combining the latter two, it is optimal to merge if and only if $T^B_{k+1}(v_H) \geq T^F_{k+1} + c_H$.

Proof. See Appendix A. □

The statement of Proposition 2 involves non-strict inequalities, with two optimal solutions at 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.

There is a certain monotonicity in the decisions.

Proposition 3. *Consider the merging problem with velocity decisions in* $\{v_L, v_H\}$. *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), in the appendix I provide a proof from first principles.

I now define the final region where it is always optimal to merge and also where optimizing driver behavior can lead to "irrational" oscillations under certain conditions.

Final Merging Zone: This is the region defined as

$$T^B_k(v_L) \geq T^B_k(v_H) \geq T^F_{k+1} + c_H = \frac{N-k+1}{\bar{v}_F} + c_H.$$

Notice that the last term is monotonically and linearly decreasing in k till the bottleneck stage N. It is always optimal to merge in this zone, and therefore, the only decision to make is whether to choose v_L or v_H.

I focus on this Final Merging Zone as it is the most interesting and critical for control purposes.

Intuitively one might expect the optimizing driver's policy to have one of the following patterns:

$$\big(\text{do not merge, high velocity}\big) \longrightarrow \big(\text{merge, high velocity}\big) \ \ldots$$

or

$$\big(\text{do not merge, high velocity}\big) \longrightarrow \big(\text{merge, low velocity}\big) \ \ldots$$

That is, try to merge at either at high or low velocities at some point and stick with that decision. But can it be optimal to oscillate between trying to merge at high velocity and then low velocity in subsequent periods? Continually? That is, if the driver just tried to merge at high velocity but failed, switch to low velocity next and try again, and if that again fails switch to high, and so on? Propositions 1 to 3 allow for this possibility.

In the remainder of the section, I show that this can indeed happen, discuss the nature of this seemingly "irrational" driving behavior, and try to gain intuition into why it happens.

3.2.1 Oscillations

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

$$\mathbb{E}_{k+1}[L] \equiv q_L \left(T^F_{k+1} + c_L\right) + (1 - q_L) T^B_{k+1}(v_L), \tag{2}$$

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 is realized), assuming that $v^e_k = v_L$; $\mathbb{E}_{k+1}[H]$ is defined similarly.

The DP recursion can be expressed in the following form:

$$T^B_k\big(v^s_k\big) = \min_{v^e_k \in \{v_L, v_H\}} \left\{ \frac{2}{v^s_k + v^e_k} + \min\left\{\mathbb{E}_{k+1}\big[v^e_k\big], T^B_{k+1}\big(v^e_k\big)\right\} \right\}. \tag{3}$$

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, which relates to the quantities that I use to characterize the structure of the optimal policy in the following way:

$$\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H] = B_k\left(v^s_k\right) - \Delta T_{k+1} - \left[\frac{2}{v^s_k + v_L} - \frac{2}{v^s_k + v_H}\right].$$

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

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

$$\frac{2}{v_L + v_L} + q_L \left(T^F_{k+1} + c_L\right) + (1 - q_L)T^B_{k+1}(v_L) \le \frac{2}{v_L + v_H} + q_H \left(T^F_{k+1} + c_H\right) + (1 - q_H)T^B_{k+1}(v_H).$$

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

It can be verified that
$$\frac{2}{v_L + v_H} - \frac{1}{v_L} \le \frac{1}{v_H} - \frac{2}{v_L + v_H} \le 0,$$
since $v_L \le v_H$. Hence, in the more interesting part of the state space, the merging zone, one can view the optimal policy as having three regions:

Region $\boldsymbol{L}$: 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 $\boldsymbol{H}$: 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 $\boldsymbol{X}$: If $\frac{2}{v_L+v_H} - \frac{1}{v_L} \le \mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H] \le \frac{1}{v_H} - \frac{2}{v_L+v_H}$, 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.

Now, to understand how 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

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

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

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

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 I explain further below.

3.2.2 Convexity and the 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. This implies the travel time at the average velocity is less than the average of the travel times at high and low velocities:

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

which can be equivalently written as

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

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:

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

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:

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

This can happen repeatedly across consecutive stages, as long as the difference of the expected values stays in region X.

3.2.3 Insight from numerical experiments

To obtain additional insight on whether the oscillations occur frequently enough to be of concern, I derive the optimal policy for a broad range of parameter values, by solving the DP (3) numerically. Figure 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:

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

Our numerical experiments also illustrate the aforementioned velocity oscillations while the driver attempts to merge. In Figure 3 I dive deeper into an example: 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.

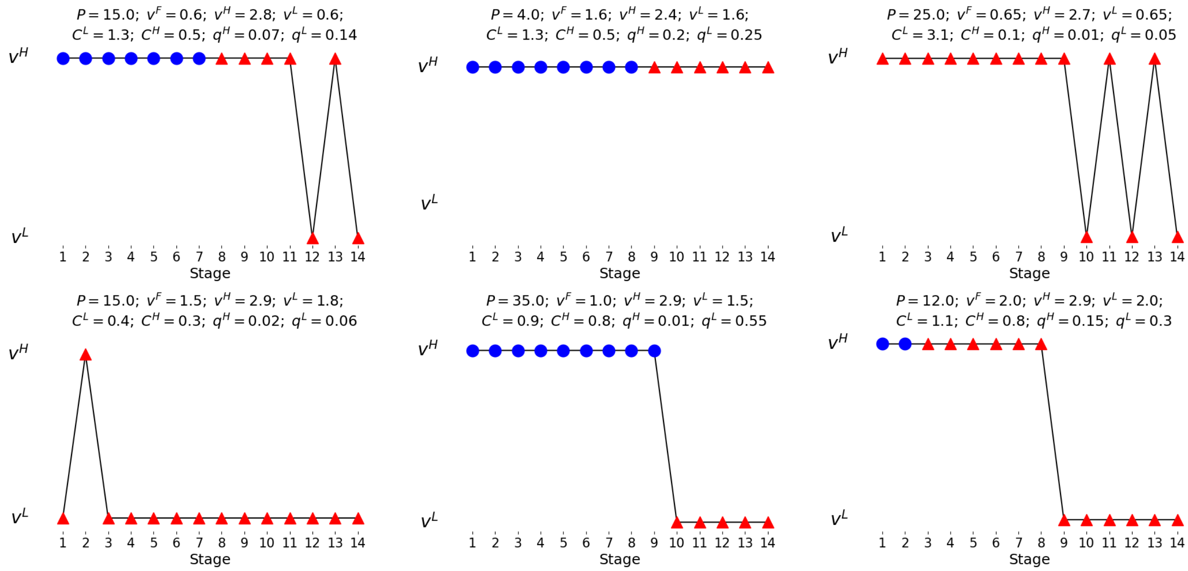

Figure 2: Examples of the optimal merging and velocity control policy for different choice of parameters P, v_H, v_L, $\bar{v}_F$, q_H, q_L, c_L and c_H, assuming $N = 15$. Optimal merging decisions are represented with the color and the shape of points (blue circle for "not merge", red triangles for "merge"), while the optimal velocities are shown along the y-axis. The corresponding stages are along the x-axis.

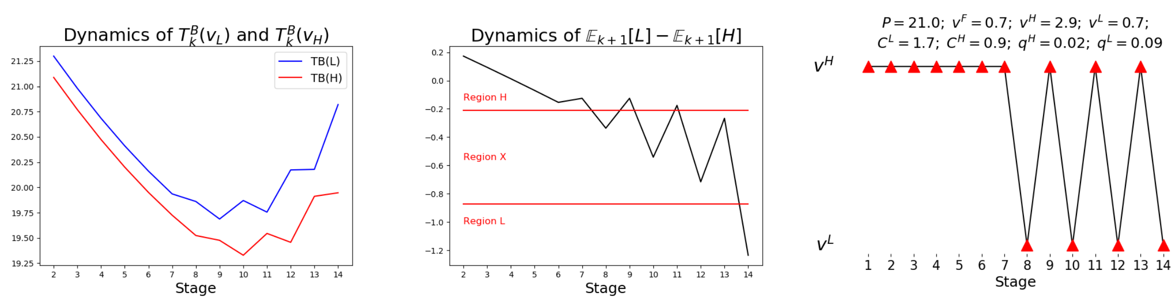

Figure 3: The dynamics of $\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H]$ lead to an optimal solution that alternates between regions H and X.

3.3 Continuous-velocity model of merging

Now, one may wonder to what extent our findings are robust to the modeling assumptions that I have made; in particular, that there are only two velocities available to the driver. To this end, I 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 $\left[v_L, v_H\right]$. The DP recursion in this case takes the following form:

$$T_k^B\left(v_k^s\right) = \min_{v_L \le v_k^e \le v_H} \left\{ \frac{2}{v_k^s + v_k^e} + \min\left\{ \underbrace{q(v_k^e)\left(T_{k+1}^F + C(v_k^e)\right) + (1 - q(v_k^e))T_{k+1}^B\left(v_k^e\right)}_{u_k=1}, \underbrace{T_{k+1}^B\left(v_k^e\right)}_{u_k=0} \right\} \right\}. \tag{4}$$

The propositions below extend in a natural way the results that I have already proved in the discrete-velocity scenario. For brevity, I omit their proofs, as they are very similar to the proofs of their counterparts.

Proposition 4. *Consider the merging problem with velocity decisions in the range $\left[v_L, v_H\right]$. 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 5. *Consider the merging problem with velocity decisions in the range $\left[v_L, v_H\right]$. 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., $u_l^* = 1,\ l = k+1, \ldots, N$.*

Eq. (4) is not amenable to further analysis, as the first-order condition for minimizing the expected travel time requires solving a fourth-order polynomial of v_k^e, with fairly complicated coefficients. Thus, I 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, I adopt a linear probability model: $q\left(v^B\right) = \alpha - \beta v^B$, where $\alpha, \beta > 0$ and $\beta v_H < \alpha \le 1$. A linear model is a good starting point for theoretical analysis in our case: Figure 14 in Online Appendix C shows the merging probability as a function of the velocity, for a frequently used model of merging. Clearly, this function can be approximated well by a straight line around its "operating point". The reader is referred to Appendix B for details. Furthermore, I use a negative-exponential function for the merging penalty: $C\left(v^B\right) = e^{\gamma - \delta v^B}$, where $\delta > 0$. Again, the reader is referred to the Online Appendix C for details.

Figure 4 presents examples of (approximately) optimal policies in the continuous-velocity case. Note that I use the same parameter sets as in Figure 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. 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, I 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 Figures 2 and 4, and three additional cases are considered.

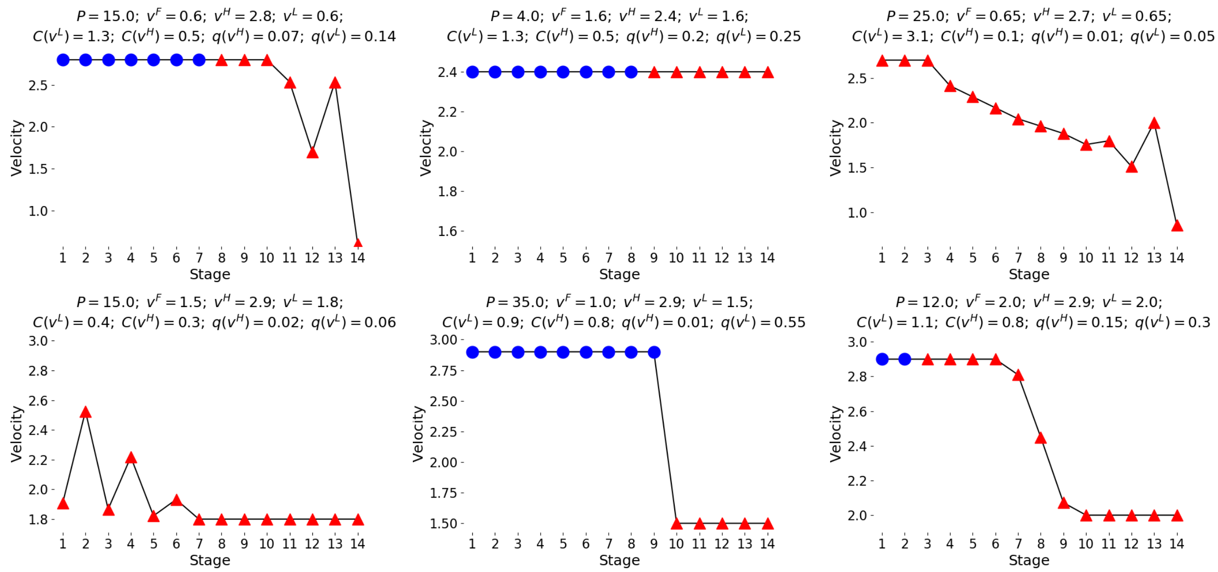

Figure 4: Examples of the optimal solution for different choice of parameters P, v_H, v_L, $\bar{v}_F$, $q(v_H)$, $q(v_L)$, $C(v_L)$ and $C(v_H)$ assuming $N = 15$. The velocity space is discretized by dividing the interval evenly into 50 discrete levels.

$\bar{v}_F$	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 case 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 this phenomenon is not an artefact of our simplifying modeling assumption, to consider just two velocities in our benchmark model, but rather an intrinsic characteristic of the problem.

4 Managing Algorithmic Drivers

Velocity oscillations while attempting to merge increase the risk of accidents. In this section I show how the traffic manager can avoid such oscillations, by setting appropriate speed limits in the Final Merging Zone.

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: $H \ldots HX \ldots XL \ldots L$. Under stronger conditions, I show that the region of oscillations can be completely avoided, so the dynamics become $H \ldots HXL \ldots L$, with at most a single X stage.

Proposition 6. *Consider the driver's merging problem with velocity decisions in $\{v_L, v_H\}$.*

1. *Assume that*
$$\frac{q_L - q_H}{\bar{v}_F} + \frac{2(1-q_L)}{v_L + v_H} - \frac{1-q_H}{v_H} + \frac{(1-q_L)(v_L - v_H)}{v_L(v_L+v_H)} \geq 0,$$
and
$$\frac{q_L - q_H}{\bar{v}_F} - \frac{2(q_L - q_H)}{v_L + v_H} + q_H q_L (c_L - c_H) + \frac{2}{v_L + v_H} - \frac{1}{v_H} + \frac{(v_L - v_H)(1-q_L)(1-q_H)}{v_L(v_L+v_H)} \geq 0.$$
Then, $\mathbb{E}_k[L] - \mathbb{E}_k[H]$ is monotonically non-increasing in k and there can be at most a single contiguous set of X regions.

2. *A stronger condition guarantees oscillations never occur. Assume that*
$$\frac{q_L - q_H}{\bar{v}_F} + \frac{2(1-q_L)}{v_L + v_H} - \frac{1-q_H}{v_H} + \frac{(1-q_L)(v_L - v_H)}{v_L(v_L+v_H)} \geq 0,$$
and
$$\frac{q_L - q_H}{\bar{v}_F} - \frac{2(q_L - q_H)}{v_L + v_H} + q_H q_L (c_L - c_H) + \frac{2}{v_L + v_H} - \frac{1}{v_H} + \frac{(v_L - v_H)(1-q_L)(1-q_H)}{v_L(v_L+v_H)} \geq \frac{(v_L - v_H)^2}{v_L v_H (v_L + v_H)}.$$
Then, the optimal solution lies in region X for, at most, one stage. Hence, the optimal policy has an intuitive "monotonic" structure:
$$(\textit{no merge, high velocity}) \longrightarrow (\textit{merge, high velocity}) \longrightarrow (\textit{merge, low velocity}),$$
which excludes oscillations.

Proof. See Appendix A. □

Corollary 1. *Assuming $\bar{v}_F \leq v_L$, 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. See Appendix A. □

Together, Proposition 1 and Corollary 1 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. What happens between these regions is where I need conditions on v_L and v_H, as a function of $\bar{v}_F$, to ensure: (a) that I do not reenter region X from H again and again; and (b) that I enter region X for at most one stage

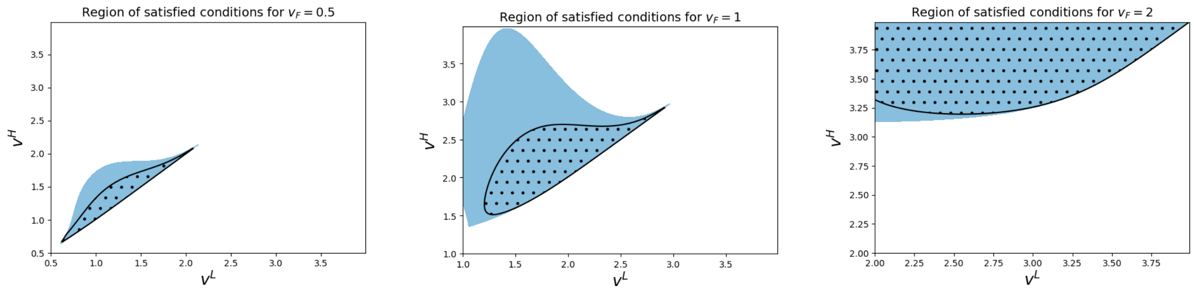

Figure 5: A region of parameters v_H, v_L, for which Proposition 6 holds, given reasonable functional forms for q_H, q_L, c_H, c_L, for different values of v_F (in normalized values; they correspond to 2.5m/s, 5m/s and 10m/s). The blue area shows the velocity region that satisfies the first part of proposition, while the dotted area the region that satisfies the second part.

before moving to region L. While valid only under our stylized model, they show that some intelligent control on speed limits is necessary to prevent these endogenous oscillations.

Interpreting the sufficient conditions on speed limits in Proposition 6 is difficult. However, as there are just two dimensions in our case, v_L and v_H, I 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. 5 I provide illustrative examples of the velocity regions that satisfy the aforementioned conditions. First, I set the values for q_L, q_H, c_L, c_H in accordance to functions of $\bar{v}_F, v_H$ and v_L that have appeared in the literature; see Online Appendix B. Then, for different values of $\bar{v}_F$, I plot the regions (v_H, v_L) where the conditions of Proposition 6 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 6 are only sufficient: in Fig. 6 I present indicative examples of optimal policies that avoid oscillations, despite not satisfying the first part of Proposition 6. 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 6, what would be the loss in efficiency, measured in aggregate travel time of the drivers? I address this question, numerically, in Fig. 7: for every such $\bar{v}_F$, I build the "safe" non-oscillating region like the ones in Fig. 5, compute the best policy within the region, and compare it to the best unconstrained policy.

As expected, the loss is near zero at around when traffic is relatively free (i.e., $\bar{v}_F \geq 8$ m/s). It also goes to zero as the conditions on the free-lane approaches complete jam (i.e., $\bar{v}_F \approx 0$ 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 Monte-Carlo Simulations

In the stylized models studied in §3, I make a number of assumptions in order to obtain concrete analytical results. It is interesting to expand the models and test the conclusions under more

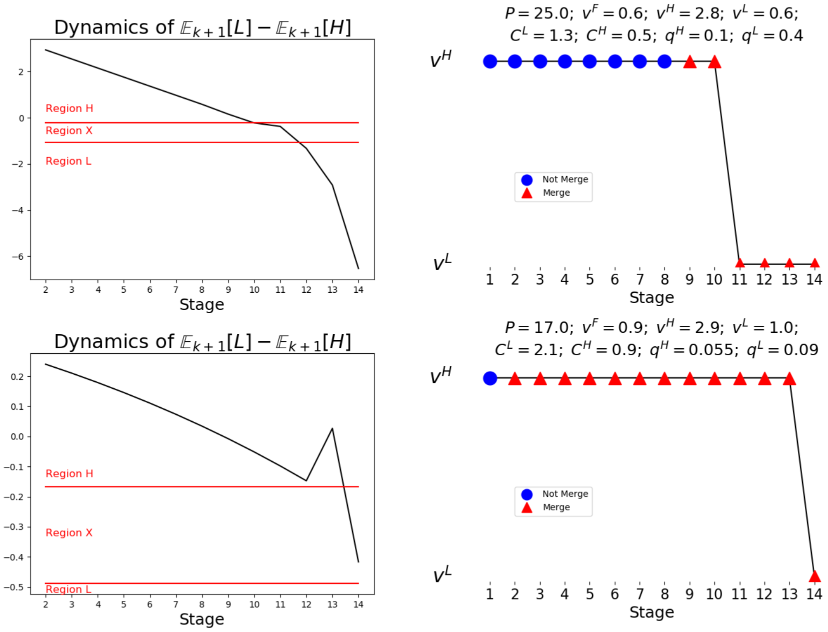

Figure 6: Two examples of DP solutions that do not oscillate despite failing to satisfy Proposition 6. Top: $\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H]$ behaves smoothly. Bottom: $\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H]$ oscillates, but outside of Region X.

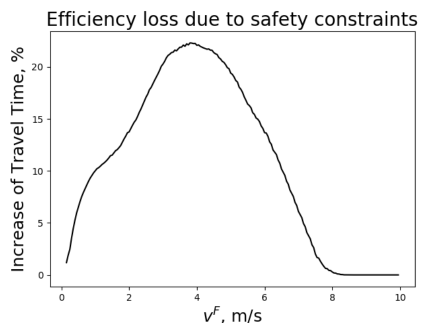

Figure 7: Additional delay, experienced by merging drivers due to imposing safety constraints, as a function of the velocity on the target lane.

realistic conditions, with multiple agents interacting with each other. Cellular Automata (CA) is a class of multi-agent simulations, with simple decision rules and their implied dynamics, employed in many fields of research where analytical solutions are hard to obtain, and widely used in traffic engineering research (Nagel et al. (1998)). In the Online Appendix D I explain in detail how I set up our discrete-event simulation framework.

5.1 Comparison to common merging guidelines

In this section I use detailed rule-based traffic simulations, modeling the behaviour of algorithm-assisted individual drivers on the free lane, to evaluate the performance of the DP-based policy in the single-car discrete-velocity case, against several merging policies recommended in practice: "early merging", "late merging", and an uncontrolled "random merging" scenario.

5.1.1 Random Merging

The first merging policy that I consider is the case of uncontrolled traffic, where drivers are not provided with any policy recommendation on when and where to merge. Every driver starts attempting to merge at a random distance from the blockage point. I consider the simplest possible case where this distance is uniformly distributed.

In the table below I summarize the performance of random merging vs. the DP approach, for different values of ρ^B and ρ^F, the density of flow on lane B and F, respectively. These values correspond to a spectrum of traffic conditions, ranging from relatively sparse to dense lane, in both cases. I denote by ΔT the percent decrease of travel time for drivers on the blocked lane under the DP policy, compared to random merging. The parameters $\bar{v}_F,\ q_H,\ q_L$, and P are estimated by averaging over 1000 CA simulations (see Online Appendix D for details on the estimation procedure), while the results for ΔT are obtained after averaging over 5000 runs.

ρ^B	ρ^F	$\bar{v}_F$	v_H	v_L	q_H	q_L	P	c_L	c_H	Optimal policy	ΔT
.03	.1	3.87	5	4	.372	.438	8.65	2.36	0.93	Merge: 100m; only v_H	3.07%
.06	.1	3.89	5	4	.375	.444	10.8	2.38	0.99	Merge: 100m; only v_H	3.68%
.09	.1	3.87	5	4	.379	.435	12.4	2.37	0.95	Merge: 120m; only v_H	2.85%
.03	.2	1.91	5	2	.139	.345	6.80	0.96	0	Merge: 60m; change to v_L at 20m	17.05%
.06	.2	1.91	5	2	.138	.340	8.99	0.98	0	Merge: 80m; change to v_L at 20m	15.38%
.09	.2	1.91	5	2	.129	.349	10.6	0.98	0	Merge: 80m; change to v_L at 20m	12.01%
.03	.3	1.21	5	2	.056	.194	6.51	0.43	0	Merge: 40m; change to v_L at 20m	26.6%
.06	.3	1.21	5	2	.051	.208	8.87	0.43	0	Merge: 40m; change to v_L at 20m	22.7%
.09	.3	1.21	5	2	.054	.205	10.1	0.43	0	Merge: 60m; change to v_L at 20m	17.6%

At higher densities of traffic on the free lane—the more interesting case for us—the difference in travel time between random merging and the DP approach becomes greater. Intuitively this makes sense, as the cost of a suboptimal decision is magnified by the congestion in the free lane, and merging closer to the blockage point becomes more appealing.

Moreover, the gap between random merging and the DP approach is quite robust with respect to the density of traffic on the blocked lane, as long as the free lane is relatively dense. I have also found the reported results to be robust to the choice of d, the length of a cell in our simulations, although I do not report on these for brevity. Together, these observations suggest that the most important primitive in terms of its effect on the performance of the DP approach is the density of traffic on the free lane.

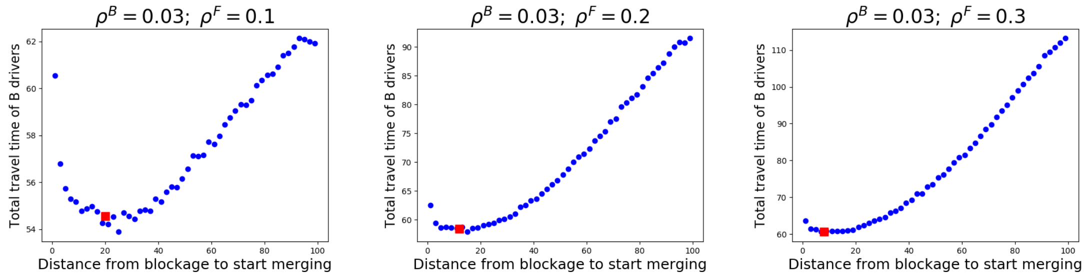

Figure 8: Comparison of the DP approach to one-threshold merging policies via Cellular-Automata simulations, for a less dense blocked lane. The merging threshold is presented on the x-axis, and the total travel time (in seconds) on the y-axis. Every point is the average of 5000 simulation runs. The red square corresponds to the proposed policy.

5.1.2 One-Threshold Merging Policies

Next, I compare the performance of the DP approach to the class of one-threshold merging policies, which includes intuitive policies such as "merge early" or "merge late". More specifically, policies included in this class recommend to the driver to start attempting to merge after a given threshold, while the velocity is always kept as high as possible.

I present our simulation results in Figures 8 and 9 (for a single driver on B and for relatively dense B-lane). Note that when the traffic on the free lane is relatively dense, the empirical performance of the DP approach is much better than the "merge early" and "merge late" policies, and very close to the best one-threshold merging policy.[7] The benefit of our framework is that it provides a principled way to achieve 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.

The results in Figures 8 and 9 also confirm one of the insights derived above, in the comparison to random merging: the (relative) performance of the DP approach is quite consistent with respect to the density of the traffic flow on the blocked lane. Correspondingly, the most critical primitive for performance evaluation purposes is the density on the free lane.

5.1.3 Bottleneck Capacity

In this section I provide further numerical comparisons of the DP approach to various heuristics. First, I compare them in terms of total travel time of B-lane drivers; second, I compare the average flow of B-lane drivers through the bottleneck; and finally, I compare both the total travel time and the flow of all drivers, including those on both the free and the blocked lane. By doing so, effectively, I compare the total capacity of the bottleneck, and demonstrate that the proposed policy does not reduce it; in fact, it often improves it. Indeed, a policy that improves traffic conditions on one lane at the expense of other lanes, would make little sense. In contrast, the proposed policy makes, on average, better use of the available resources.

[7]As expected, an optimal solution to the single-car problem does not always have the best performance, in simulations, among threshold policies. This could be due to imperfect estimation of the model parameters, or the fact that our model does not account for the influence of merging drivers on the free lane. Typically, the DP approach recommends to merge slightly closer to the blockage than the (empirically) best threshold policy, which likely means that either the merging probability is overestimated, or the late merging penalty P is underestimated, or both. However, since the performance of the DP policy for a single car on the blocked lane is even closer to the optimum, the most of the gap is likely to be attributed to unaccounted interactions of the blocked-lane cars.

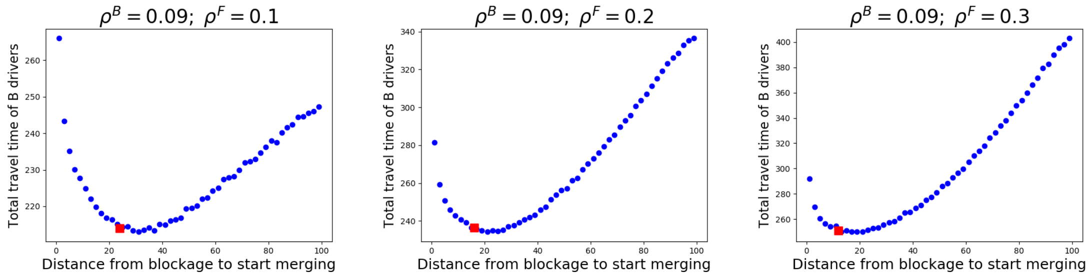

Figure 9: Comparison of the DP approach to one-threshold merging policies via Cellular-Automata simulations, for a more dense blocked lane. The merging threshold is presented on the x-axis, and the total travel time (in seconds) on the y-axis. Every point is the average of 5000 simulation runs. The red square corresponds to the proposed policy.

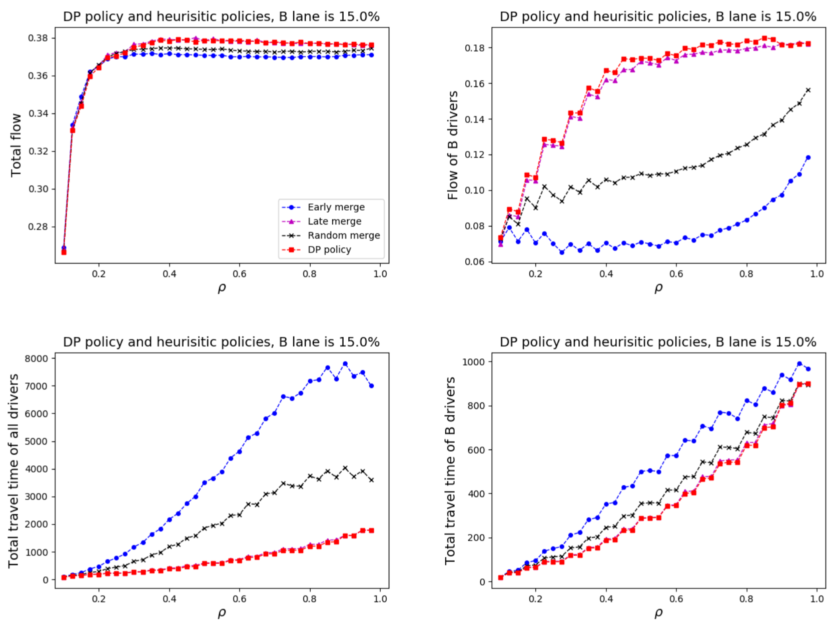

Figure 10: Comparison of the DP approach to three other popular heuristic policies, for different traffic densities. Traffic on B-lane is light (15% of total density)

In Figures 10, 11, 12 and 13 I plot the aforementioned measures for the different merging policies, for four different levels of traffic intensity on B-lane: 15%, 30%, 45% and 60%. Note that, in general, increasing the density leads to a steady increase of total travel time. The traffic flow metric behaves differently: for the pooled metric, it grows fast until the bottleneck capacity, and then almost stabilizes. Different policies can lead to slightly different bottleneck capacity. Also, the flow of B-lane drivers can behave differently, for example, it can decrease at very high density, due to the fact that it become harder for B-lane drivers to leave the bottleneck when there are many F-lane drivers. By and large though, the results supports our main theses: the DP provides a principled way to achieve near-optimal performance under any traffic conditions, whereas the different heuristic policies may or may not perform well depending on the circumstances.

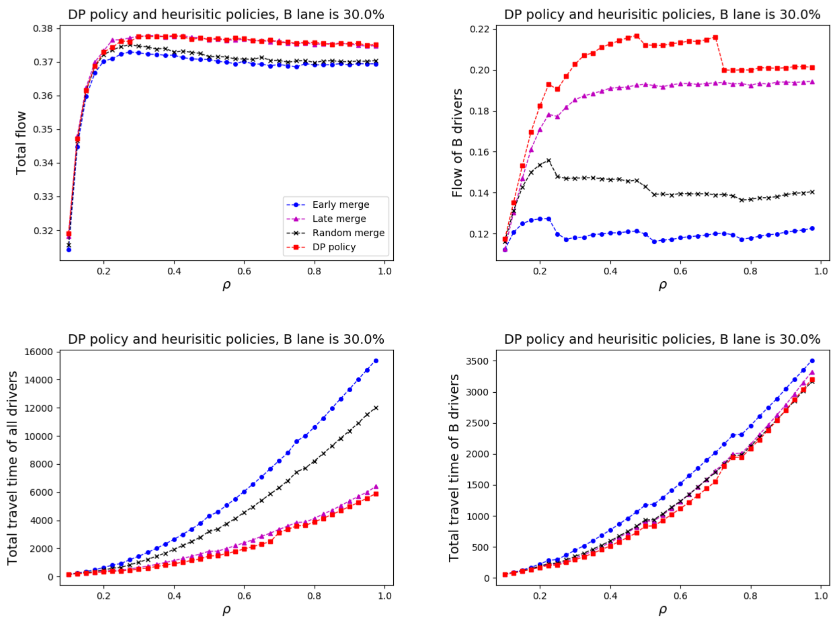

Figure 11: Comparison of the DP approach to three other popular heuristic policies, for different traffic densities. Traffic on B-lane is significant (30% of total density)

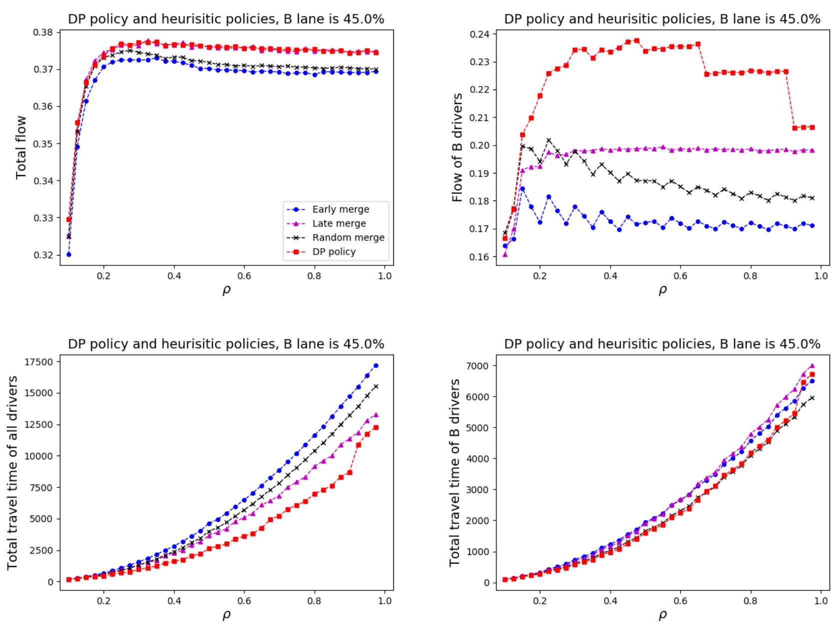

Figure 12: Comparison of the DP approach to three other popular heuristic policies, for different traffic densities. Traffic is split almost evenly (45% of total density on B-lane)

6 Conclusions

Algorithmic-driving technologies and vehicle-to-vehicle communication are reaching maturity, and bring urgency to research on driver behavior, incentives, and control policies for "optimal"

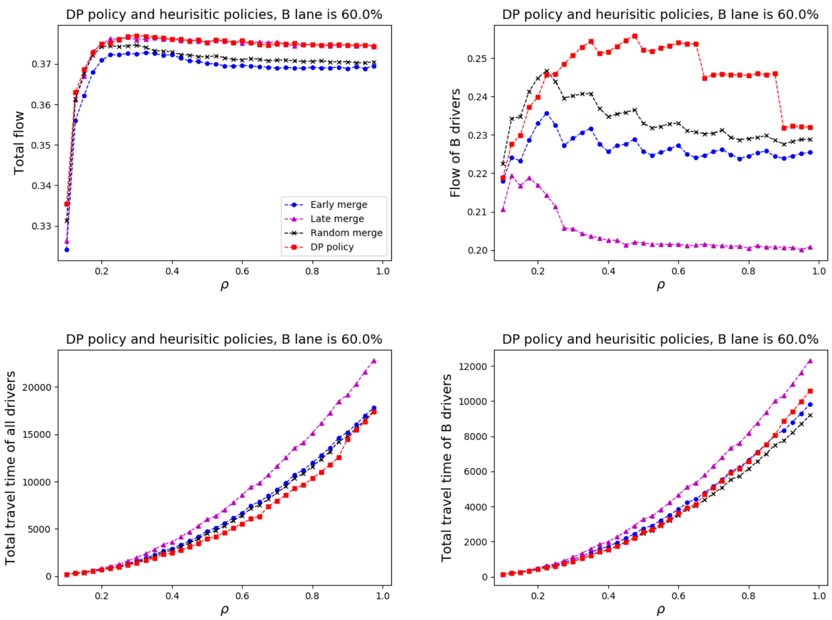

Figure 13: Comparison of the DP approach to three other popular heuristic policies, for different traffic densities. Traffic on B-lane is heavy (60% of total density on B-lane)

traffic flow. And while decades of traffic studies have built elaborate and complicated theories on the macroscopic phenomena observed on highways, we still lack a deeper understanding of the underlying causes for fluctuations, which makes it difficult to devise good intervention and control policies at the individual-driver level.

To this end, I 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 behavioural 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. I validate our theoretical analysis via extensive discrete-event simulations under real-life scenarios, where I evaluate the macroscopic impact of the DP policy against various merging heuristics.

Traffic modeling and analysis is extremely challenging as multiple agents interact in a dynamically changing environment. However, algorithm-assisted drivers are somewhat easier to model at the microscopic level and the analysis more tractable potentially leading to actionable insights, and I believe this represents a new and exciting area of research, with a huge potential for making an impact on our daily lives.

As the main direction for future research I consider a more rigorous treatment of the case with multiple cars on the blocked lane and the analysis of the continuous velocity case. There is a vast array of interesting questions to consider when one combines this with free-lane driver behavior, and even more intriguing ones when drivers are modeled as acting strategically.

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

Acknowledgements

This work was supported by the ECO2013-41131-R and ECO2016-75905-R (AEI/FEDER, UE) grants from the Spanish Ministry of Economy, Industry and Competitiveness.

References

Acemoglu, D, R Johari, A Ozdaglar. 2007. Partially optimal routing. *IEEE Journal on Selected Areas in Communications* **25**(6) 1–13.

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

Agarwal, S., P. Kachroo, S. Contreras, S. Sastry. 2015. Feedback-coordinated ramp control of consecutive on-ramps using distributed modeling and godunov-based satisfiable allocation. *IEEE Transactions on Intelligent Transportation Systems* **16**(5) 2384–2392.

Ahmed, K. I. 1999. Modeling drivers' acceleration and lane changing behavior. Ph.D. thesis, MIT, Cambridge, MA.

Ahn, Soyoung, Michael J Cassidy. 2007. Freeway traffic oscillations and vehicle lane-change maneuvers. Transportation and Traffic Theory 2007.

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

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

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

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

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

Baykal-Gürsoy, M., W. Xiao, K. Ozbay. 2009. Modeling traffic flow interrupted by incidents. *European Journal of Operational Research* **195**(1) 127–138.

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

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

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

Choudhury, Charisma, Moshe E. Ben-akiva, Tomer Toledo, Anita Rao, Gunwoo Lee. 2007. State dependence in lane changing. *Transportation and Traffic Theory* 711–734.

Choudhury, Charisma F., Moshe E. Ben-Akiva. 2013. Modelling driving decisions: a latent plan approach. *Transportmetrica A: Transport Science* **9**(6) 546–566.

Chowdhury, Debashish, Ludger Santen, Andreas Schadschneider. 2000. Statistical physics of vehicular traffic and some related systems. *Physics Reports* **329**(4) 199 – 329.

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

Correa, J, A Schultz, N Stier-Moses. 2004. Selfish routing in capacitated networks. *Mathematics of Operations Research* **29**(4) 961–976.

Daganzo, C.F. 2002a. A behavioral theory of multi-lane traffic flow part i: Long homogeneous freeway sections. *Transportation Research B* **36** 131–158.

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

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

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

DOT, US. 2018. Preparing for the future of transportation: Automated vehicles 3.0. Tech. rep., Department of Transportation.

Duranton, G., M.A. Turner. 2011. The fundamental law of road congestion: Evidence from us cities. *The American Economic Review* **101**(6) 2616–2652.

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

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

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

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

Fritzsche, H-T. 1994. A model for traffic simulation. *Transportation Engininering Contribution* **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 Research Board Proceedings* **14** 448–477.

Gregoire, J, X Qian, E Frazzoli, A de la Fortelle, T Wongpiromsarn. 2015. Capacity-aware backpressure signal control. *IEEE Transactions on Control of Networked Systems* **2**(2) 164–173.

Grillo, L., T. Datta, C. Hartner. 2008. Dynamic late lane merge system at freeway construction work zones. *Transportation Research Record: Journal of the Transportation Research Board* **2055** 3–10.

Han, Y.-S., S.-K. Ko. 2012. Analysis of a cellular automaton model for car traffic with a junction. *Theoretical Computer Science* **450** 54–67.

Hawkes, A.G. 1968. Gap acceptance in road traffic. *Journal of Applied Probability* **5** 84–92.

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

Helbing, Dirk. 2001. Traffic and related self-driven many-particle systems. *Rev. Mod. Phys.* **73** 1067–1141.

Hidas, P. 2002. Modelling lane changing and merging in microscopic traffic simulation. *Transportation Research Part C: Emerging Technologies* **10** 351–371.

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

Hounsell, N.B., S.R. Barnard, M. McDonald, D. Owens. 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., C. Roncoli, I. Papamichail, M. Papageorgiou. 2015. Feedback-based mainstream traffic flow control for multiple bottlenecks on motorways. *IEEE Transactions on Intelligent Transportation Systems* **16**(2) 610–621.

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

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

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

Kesting, Arne, Martin Treiber, Dirk Helbing. 2007. General lane-changing model mobil for car-following models. *Transportation Research Record: Journal of the Transportation Research Board* **1999** 86–94.

Kita, H. 1999. A merging-giveway interaction model of cars in a merging section: a game theoretic analysis. *Transportation Research Part A: Policy and Practice* **33**(3-4) 305–312.

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

Kotsialos, A., M. Papageorgiou. 2004. Nonlinear optimal control applied to coordinated ramp metering. *IEEE Transactions on Control Systems Technology* **12**(6) 920–933.

Lago, A., C. Daganzo. 2007. Spillovers, merging traffic and the morning commute. *Transportation Research Part B: Methodological* **41**(6) 670–683.

Laval, Jorge A., Carlos F. Daganzo. 2006. Lane-changing in traffic streams. *Transportation Research Part B: Methodological* **40**(3) 251–264.

Le, T, P Kovacs, N Walton, H Vu, L Andrew, S Hoogendorn. 2015. Decentralized signal control for urban road networks. *Transportation Research Part 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., G. Hyman. 2012. Modelling motorway merge: the current practice in the UK and towards establishing general principles. *Transport Policy* **24** 199–210.

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

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

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

McQueen, J., R.G. Miller Jr. 1960. Optimal persistence policies. *Operations Research* **8**(3) 362–380.

Miller, A.J. 1961. A queueing model for road traffic flow. *Journal of the Royal Statistical Society. Series B (Methodological)* **23**(1) 64–90.

Nagel, K., M. Schreckenberg. 1992. A cellular automaton model for freeway traffic. *Journal de Physique I* **2**(12).

Nagel, K., D.E. Wolf, P. Wagner, P. Simon. 1998. Two-lane traffic rules for cellular automata: A systematic approach. *Physical Review E* **58**(2) 1425–1437.

New York Times. 2016. Cars talking to one another? They could under proposed safety rules. URL https://nyti.ms/2jQkpt4.

New York Times. 2017. Self-driving cars can't cure traffic, but economics can. URL https://nyti.ms/2mEvPlk.

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

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

Rahman, Mizanur, Mashrur Chowdhury, Yuanchang Xie, Yiming He. 2013. Review of microscopic lane-changing models and future research opportunities. *IEEE Transactions on Intelligent Transportation Systems* **14**(4) 1942–1956.

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

Rickert, M., K. Nagel, M. Schreckenberg, A. Latour. 1996. Two lane traffic simulations using cellular automata. *Physica A: Statistical Mechanics and its Applications* **231**(4) 534–550.

Rios-Torres, J., A.A. Malikopoulos. 2017a. Automated and cooperative vehicle merging at highway on-ramps. *IEEE Transactions on Intelligent Transportation Systems* **18**(4) 780–789.

Rios-Torres, J., A.A. Malikopoulos. 2017b. A survey on the coordination of connected and automated vehicles at intersections and merging at highway on-ramps. *IEEE Transactions on Intelligent Transportation Systems* **18**(5) 1066–1077.

Shalev-Shwartz, Shai, Shaked Shammah, Amnon Shashua. 2016. Safe, multi-agent reinforcement learning for autonomous driving. *arXiv* **1610.03295**. URL http://arxiv.org/abs/1610.03295.

Shalev-Shwartz, Shai, Shaked Shammah, Amnon Shashua. 2017. On a formal model of safe and scalable self-driving cars. *CoRR* **abs/1708.06374**. URL url={http://arxiv.org/abs/1708.06374}.

Sheu, Jiuh-Biing. 2013. Microscopic traffic behaviour modelling and simulation for lane-blocking arterial incidents. *Transportmetrica A: Transport Science* **9**(4) 335–357.

Sun, Daniel, Lily Elefteriadou. 2010. Research and implementation of lane-changing model based on driver behavior. *Transportation Research Record: Journal of the Transportation Research Board* **2161**.

Tarko, A.P., D. Shamo, J. Wasson. 1999. Indiana lane merge system for work zones on rural freeways. *Journal of Transportation Engineering* **125**(5) 415–420.

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

Toledo, Tomer, Haris N. Koutsopoulos, Moshe E. Ben-Akiva. 2003. Modeling integrated lane-changing behavior. *Transportation Research Record* **1857**(1) 30–38.

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. *American Economic Review* **59** 251–261.

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

Wardrop, J.G. 1952. Some theoretical aspects of road traffic research. *Proceedings of the Institution of Civil Engineers* **1**(3) 325–362.

Wei, Heng, Eric Meyer, Joe Lee, Chuen Feng. 2000. Characterizing and modeling observed lane-changing behavior: Lane-vehicle-based microscopic simulation on urban street network. *Transportation Research Record* **1710**(1) 104–113.

Weiss, G.H., A.A. Maradudin. 1962. Some problems in traffic delay. *Operations Research* **10** 74–104.

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

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

Zheng, Zuduo. 2014. Recent developments and research needs in modeling lane changing. *Transportation Research Part B: Methodological* **60** 16–32.

Zheng, Zuduo, Soyoung Ahn, Danjue Chen, Jorge Laval. 2013. The effects of lane-changing on the immediate follower: Anticipation, relaxation, and change in driver characteristics. *Transportation Research Part C: Emerging Technologies* **26** 367–379.

Zheng, Zuduo, Soyoung Ahn, Christopher M. Monsere. 2010. Impact of traffic oscillations on freeway crash occurrences. *Accident Analysis & Prevention* **42**(2) 626–636.

Appendices

A Proofs of Theoretical Results

A.1 Technical Lemmas

I 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 beginning 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

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

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 beginning of the stage is v_L, and working similarly. □

A.2 Proof of Proposition 2

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

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

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, I have that $T_{k+1}^B(v_H) \leq T_{k+1}^F + c_H$, since $q_H > 0$.

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

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

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, I 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), I 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

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

if I 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, I 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.

I 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 I do not use strict inequalities in the statement or the proof of the result.

A.3 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, $\bar{v}_F$, therefore

$$T_{k+1}^F = \frac{1}{\bar{v}_F} + T_{k+2}^F \implies T_{k+1}^F + c_H = \frac{1}{\bar{v}_F} + T_{k+2}^F + c_H \implies T_{k+1}^B(v_H) \geq \frac{1}{\bar{v}_F} + T_{k+2}^F + c_H, \quad (5)$$

since Proposition 2 implies that $T^B_{k+1}(v_H) \geq T^F_{k+1} + c_H$. Now, Eq (1), with $u^*_k = 1$, implies that

$$\begin{aligned}
T^B_{k+1}(v_H) &= \min_{v^e_{k+1} \in \{v_L, v_H\}} \left\{ \frac{2}{v_H + v^e_{k+1}} + \left\{ q(v^e_{k+1})(T^F_{k+2} + C(v^e_k)) + (1 - q(v^e_{k+1}))T^B_{k+2}(v^e_{k+1}) \right\} \right\} \\
&\leq \frac{1}{v_H} + q_H \left(T^F_{k+2} + c_H\right) + (1 - q_H)T^B_{k+2}(v_H).
\end{aligned}$$

Combined with Eq. (5), I have that

$$\begin{aligned}
& \frac{1}{\bar{v}_F} + T^F_{k+2} + c_H \leq \frac{1}{v_H} + q_H \left(T^F_{k+2} + c_H\right) + (1 - q_H)T^B_{k+2}(v_H) \\
\iff \quad & (1 - q_H)\left(T^F_{k+2} + c_H\right) \leq \frac{1}{v_H} - \frac{1}{\bar{v}_F} + (1 - q_H)T^B_{k+2}(v_H) \\
\iff \quad & T^F_{k+2} + c_H \leq T^B_{k+2}(v_H),
\end{aligned}$$

since $\bar{v}_F \leq v_H$. This is equivalent to $u^*_{k+1} = 1$, according to Proposition 2.

A.4 Proof of Propositions 6 and Corollary 1

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, I use the fact that

$$\Delta E_k = \frac{q_L - q_H}{\bar{v}_F} + (1 - q_L)\left(T^B_k(v_L) - T^B_{k+1}(v_L)\right) - (1 - q_H)\left(T^B_k(v_H) - T^B_{k+1}(v_H)\right). \tag{6}$$

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 $v^{*e}_k = v_L$, irrespective of the velocity at the beginning of stage k. Eq. (6) implies that

$$\begin{aligned}
\Delta E_k &= \frac{q_L - q_H}{\bar{v}_F} + (1 - q_L)\left(\frac{1}{v_L} + q_L(T^F_{k+1} + c_L) + (1 - q_L)T^B_{k+1}(v_L) - T^B_{k+1}(v_L)\right) \\
&- (1 - q_H)\left(\frac{2}{v_H + v_L} + q_L(T^F_{k+1} + c_L) + (1 - q_L)T^B_{k+1}(v_L) - T^B_{k+1}(v_H)\right) \\
&= \frac{q_L - q_H}{\bar{v}_F} + \frac{1 - q_L}{v_L} - \frac{2(1 - q_H)}{v_L + v_H} - (1 - q_L)q_L T^B_{k+1}(v_L) - (1 - q_H)(1 - q_L)T^B_{k+1}(v_L) \\
&+ (1 - q_H)T^B_{k+1}(v_H) + (q_H - q_L)q_L\left(T^F_{k+1} + c_L\right) \\
&= \frac{q_L - q_H}{\bar{v}_F} + \frac{1 - q_L}{v_L} - \frac{2(1 - q_H)}{v_L + v_H} + (1 - q_H)\left(T^B_{k+1}(v_H) - T^B_{k+1}(v_L)\right) \\
&+ (q_H - q_L)q_L\left(T^F_{k+1} + c_L - T^B_{k+1}(v_L)\right).
\end{aligned}$$

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

$$\frac{2}{v^s_k + v_L} + q_L(T^F_{k+1} + c_L) + (1 - q_L)T^B_{k+1}(v_L) \leq \frac{2}{v^s_k + v_L} + T^B_{k+1}(v_L).$$

Rearranging terms, I have that $T^F_{k+1} + c_L \leq T^B_{k+1}(v_L)$, 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 $T^B_{k+1}(v_H) - T^B_{k+1}(v_L)$ 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 establishes Corollary 1 also.

Region X: If $\mathbb{E}_{k+1}[L] - \mathbb{E}_{k+1}[H]$ is in region X, then it is optimal to merge and $v^{*e}_k \neq v^s_k$. Eq. (6) implies that

$$\begin{aligned}
\Delta E_k = & \frac{q_L - q_H}{\bar{v}_F} + (1-q_L)\left(\frac{2}{v_H + v_L} + q_H(T^F_{k+1} + c_H) + (1-q_H)T^B_{k+1}(v_H) - T^B_{k+1}(v_L)\right) \\
& - (1-q_H)\left(\frac{2}{v_H + v_L} + q_L(T^F_{k+1} + c_L) + (1-q_L)T^B_{k+1}(v_L) - T^B_{k+1}(v_H)\right) \\
= & \frac{q_L - q_H}{\bar{v}_F} - \frac{2(q_L - q_H)}{v_L + v_H} + (q_H - q_L)T^F_{k+1} + (1-q_L)q_H c_H - (1-q_H)q_L c_L \\
& + (1-q_L)(1-q_H)\big(T^B_{k+1}(v_H) - T^B_{k+1}(v_L)\big) + (1-q_H)T^B_{k+1}(v_H) - (1-q_L)T^B_{k+1}(v_L).
\end{aligned}$$

By adding and subtracting $q_L T^B_{k+1}(v_H)$, I can rewrite the expression as follows:

$$\begin{aligned}
\Delta E_k = & \frac{q_L - q_H}{\bar{v}_F} - \frac{2(q_L - q_H)}{v_L + v_H} + (q_H - q_L)T^F_{k+1} + (1-q_L)q_H c_H - (1-q_H)q_L c_L \\
& + (1-q_L)(1-q_H)\big(T^B_{k+1}(v_H) - T^B_{k+1}(v_L)\big) + (1-q_L)T^B_{k+1}(v_H) \\
& - (1-q_L)T^B_{k+1}(v_L) + (q_L - q_H)T^B_{k+1}(v_H) \\
= & \frac{q_L - q_H}{\bar{v}_F} - \frac{2(q_L - q_H)}{v_L + v_H} + (1-q_L)q_H c_H - (1-q_H)q_L c_L \\
& + (1-q_L)(2-q_H)\big(T^B_{k+1}(v_H) - T^B_{k+1}(v_L)\big) + (q_L - q_H)\big(T^B_{k+1}(v_H) - T^F_{k+1}\big).
\end{aligned} \tag{7}$$

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

$$\begin{aligned}
& E_{k+1}[L] - E_{k+1}[H] \\
= & (q_L - q_H)T^F_{k+1} + q_L c_L - q_H c_H + (1-q_L)T^B_{k+1}(v_L) - (1-q_H)T^B_{k+1}(v_H) \\
= & (q_L - q_H)T^F_{k+1} + q_L c_L - q_H c_H + (1-q_L)T^B_{k+1}(v_L) - (1-q_L)T^B_{k+1}(v_H) - (q_L - q_H)T^B_{k+1}(v_H) \\
= & (q_L - q_H)\big(T^F_{k+1} - T^B_{k+1}(v_H)\big) + q_L c_L - q_H c_H - (1-q_L)\big(T^B_{k+1}(v_H) - T^B_{k+1}(v_L)\big),
\end{aligned}$$

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

$$(q_L - q_H)(T^B_{k+1}(v_H) - T^F_{k+1}) \geq \frac{2}{v_L + v_H} - \frac{1}{v_H} + q_L c_L - q_H c_H - (1-q_L)\big(T^B_{k+1}(v_H) - T^B_{k+1}(v_L)\big).$$

Combining the above lower bound with Eq. (7), I obtain:

$$\Delta E_k \geq \frac{q_L - q_H}{\bar{v}_F} - \frac{2(q_L - q_H)}{v_L + v_H} + q_H q_L (c_L - c_H) + \frac{2}{v_L + v_H} - \frac{1}{v_H} + (1 - q_L)(1 - q_H)\big(T^B_{k+1}(v_H) - T^B_{k+1}(v_L)\big).$$

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

$$\Delta E_k \geq \frac{q_L - q_H}{\bar{v}_F} - \frac{2(q_L - q_H)}{v_L + v_H} + q_H q_L (c_L - c_H) + \frac{2}{v_L + v_H} - \frac{1}{v_H} + \frac{(v_L - v_H)(1 - q_L)(1 - q_H)}{v_L(v_L + v_H)}. \quad (8)$$

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 beginning of stage k. Eq. (6) implies that

$$\begin{aligned}
\Delta E_k &= \frac{q_L - q_H}{\bar{v}_F} + (1 - q_L)\left(\frac{2}{v_H + v_L} + q_H(T^F_{k+1} + c_H) + (1 - q_H)T^B_{k+1}(v_H) - T^B_{k+1}(v_L)\right) \quad (9)\\
&- (1 - q_H)\left(\frac{1}{v_H} + q_H(T^F_{k+1} + c_H) + (1 - q_H)T^B_{k+1}(v_H) - T^B_{k+1}(v_H)\right)\\
&= \frac{q_L - q_H}{\bar{v}_F} + \frac{2(1 - q_L)}{v_L + v_H} - \frac{1 - q_H}{v_H} + q_H(q_L - q_H)\left(T^B_{k+1}(v_H) - T^F_{k+1} - c_H\right)\\
&+ (1 - q_L)\left(T^B_{k+1}(v_H) - T^B_{k+1}(v_L)\right)\\
&\geq \frac{q_L - q_H}{\bar{v}_F} + \frac{2(1 - q_L)}{v_H + v_L} - \frac{1 - q_H}{v_H} + \frac{(1 - q_L)(v_L - v_H)}{v_L(v_L + v_H)},
\end{aligned}$$

where, to derive the last inequality, I 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 (8) and (9) hold, which is precisely the requirement of Part (i) of Proposition 6, then ΔE_k is nonnegative for every k, hence $\mathbb{E}_k[L] - \mathbb{E}_k[H]$ is monotonically nonincreasing in k;

The proof of the second part of Proposition 6 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. (8) with the fact that the length of the region X is equal to

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

I derive the following sufficient condition:

$$\frac{q_L - q_H}{\bar{v}_F} - \frac{2(q_L - q_H)}{v_L + v_H} + q_H q_L (c_L - c_H) + \frac{2}{v_L + v_H} - \frac{1}{v_H} + \frac{(v_L - v_H)(1 - q_L)(1 - q_H)}{v_L(v_L + v_H)} \geq \frac{(v_L - v_H)^2}{v_L v_H (v_L + v_H)} \quad (10)$$

Summarizing, if both conditions (9) and (10) hold, which is precisely the requirement of Part (ii) of Proposition 6, then the solution lies in region X at most once, hence the optimal policy exhibits no oscillations.

Online Appendices

B Numerical Experiment Setup

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, a more plausible/concrete model of merging is necessary. It is well known that the probability of merging depends on the density in the free lane, d^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 I review the existing literature in order to develop functional forms $q(v^B, \bar{v}_F)$ and $P(\bar{v}_F)$, where v^B represents the velocity in the blocked lane, which I 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, I define the merging probability as follows:

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

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

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

$$G_c^{lead} = \exp(\alpha_0 + \frac{\alpha_1}{1+\exp(-\max\{0, \Delta V^{lead}\})} + \alpha_2 \min\left\{0, \Delta V^{lead}\right\} + \frac{\alpha_3 d_{nt}}{1+\exp(\alpha_4+\alpha_5\nu)} + \alpha_6\nu + \epsilon^{lead}), \tag{11}$$

and

$$G_c^{lag} = \exp(\beta_0 + \beta_1 \max\left\{0, \Delta V^{lag}\right\} + \beta_2 \min\left\{0, \Delta V^{lead}\right\} + \frac{\beta_3 d_{nt}}{1+\exp(\beta_4+\beta_5\nu)} + \beta_6\nu + \beta_7 \max\left\{0, a^{lag}\right\} + \epsilon^{lag}), \tag{12}$$

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.

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

$$G_c = l + G_c^{lead} + G_c^{lag},$$

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, I also make the following assumptions:

i. I omit the idiosyncratic aggressiveness term ν, due to the fact that I only have one driver on lane B;

ii. I set $d_{nt} = 0$, suggesting that the critical gap does not depend on the distance from the blockage point in our model;

iii. I 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 I set them to zero as an approximation.

Summarizing, G_c is determined by the following equation:

$$G_c(v^B, \bar{v}_F) = 4 + \exp(0.627 + \frac{1.90}{1 + \exp(-\max\{0, \bar{v}_F - v^B\})} - 0.314 \min\{0, \bar{v}_F - v^B\}) \\ + \exp(0.509 + 0.116 \max\{0, \bar{v}_F - v^B\} + 0.034 \min\{0, \bar{v}_F - v^B\}), \quad (13)$$

where I 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, I first need to estimate the mean density on the free lane, which is typically a function of the velocity, i.e., $d^F(\bar{v}_F)$. 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). I use one of the popular estimated functional forms (see Del Castillo and Benítez (1995)). If I adjust the units of the function to meters and meters per second, then the function takes the following form:

$$d^F(\bar{v}_F) = 7.48 \exp(\frac{\bar{v}_F}{8.05}) - l, \quad (14)$$

where, again, I let $l = 4$m. (I 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 $\bar{v}_F$, this is equivalent to modeling car arrivals to the studied road section as a Poisson process. Hence, in our numerical experiments I assume that the available gap is an exponential random variable with parameter $1/d^F(\bar{v}_F)$.

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

$$q(v^B, \bar{v}_F) = \exp(-\frac{G_c(v^B, \bar{v}_F)}{d^F(\bar{v}_F)}),$$

where $G_c(v^B, \bar{v}_F)$ and $d^F(\bar{v}_F)$ are given by Equations (13) and (14), respectively.

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

$$P(\bar{v}_F) = m + \frac{1}{\bar{v}_F} \cdot \frac{1}{q(0, \bar{v}_F)} + C(0). \quad (15)$$

I 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: $1/\bar{v}_F$ is the average time required to "observe" a new gap, and $1/q(0, \bar{v}_F)$ the average number of attempts needed to find a suitable gap (the mean of corresponding geometric distribution). In our numerical examples I use $m = 2$,

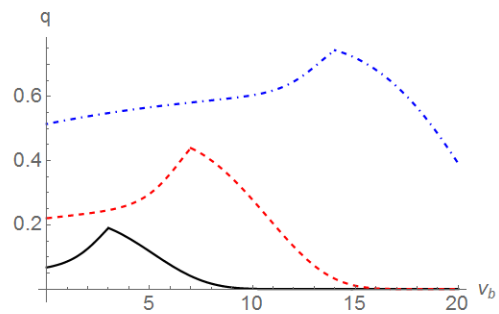

Figure 14: The function $q(v^B, \bar{v}_F)$ for $\bar{v}_F = 3$m/s (black solid line), $\bar{v}_F = 7$m/s (red dashed line), and $\bar{v}_F = 14$m/s (blue dot-dashed line).

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. The concrete models that can be used for $C(v^B)$ are discussed below.

In the blocked-lane velocity optimization problem of §4, before I proceed to any numerical solution, I also have to determine appropriate values for d, N, and v^{max}. In our theoretical analysis, I normalize d to 1, but for the numerical experiments I need an actual value since, in turn, it determines the values of the normalized velocities in the model. In general, d relates to our belief about how often the driver checks for new merging opportunities, and it can vary depending on the behavior of the driver or the traffic conditions. I use $d = 5$m as a default value. The choice of the number of stages, N, is closely related to the choice of d, as together they determine the length of the road section which I am considering. For $d = 5$, a value of N around $15 - 20$ seems reasonable. Regarding the maximum allowed speed, I use $v^{max} = 20$m/s in our numerical examples, as this is the limit in the estimated functional form that I employ.

C 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 manoeuver. This effect is called "relaxation" and may be an important factor for an optimal merging decision.

I 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 $\bar{v}_F$ as $C(g, v^B, \bar{v}_F)$. As before, I count only the pooled gap, rather than discriminating between lead gap and lag gap. Then, the (expected) merging penalty that I use in our DP can essentially be written as follows:

$$C(v^B) = \mathbb{E}[C(g, v^B, \bar{v}_F)|g > G_c(v^B, \bar{v}_F)].$$

To show analytically the required characterization, I 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 $\bar{v}_F$ affect the expected merging penalty only through critical gap $G_c(v^B, \bar{v}_F)$, and therefore the merging penalty $C(g, v^B, \bar{v}_F) = C(g)$ depends only on g;

iii. For $g \geq 0$ the function $C(g)$ maps to the non-negative values and is strictly decreasing. I show results a for negative-exponential function $e^{\gamma-\delta g}$, for parameters $\gamma \geq 0, \delta > 0$. I briefly discuss below also an alternative model with function $e^{\gamma-\delta(g-G_c(v^B,\bar{v}_F))}$. These two functions leads to clean analytical results. Qualitative results for reciprocal function $1/(\gamma+\delta g)$ are similar, although less tractable.

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

$$C(v^B) = \mathbb{E}[C(g, v^B, \bar{v}_F)|g > G_c] = \frac{1}{1-F(G_c)}\int_{G_c}^{\infty} f(g)C(g) = \frac{1}{e^{-\lambda G_c}}\int_{G_c}^{\infty} \lambda e^{\gamma-(\lambda+\delta)g} = \\ = \frac{\lambda e^{\gamma-\delta G_c}}{\lambda+\delta}. \quad (16)$$

Since it is known (in particular, it follows from the models I discussed above) that $G_c(v_H, \bar{v}_F) > G_c(v_L, \bar{v}_F)$, 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 I would consider a similar functional form for $C(g, v^B, \bar{v}_F)$, but having the penalty depending on the difference between the actual gap and some constant minimally required gap $g_0 < G_c(v^B, \bar{v}_F)$. Indeed, this leads to the result:

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

which is the same model with a higher constant.

I 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. I can conclude, that under reasonable modelling assumptions, ex-post merging penalty in the optimal control and velocity problem is such that $C(v_H) \leq C(v_L)$ for $v_H > v_L$.

I 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 $(d^F(\bar{v}_F) - g/2)/\bar{v}_F$, that is the time to cover the difference between the gap and the normal for that velocity inter-car distance, at velocity $\bar{v}_F$ (in case $d^F(\bar{v}_F) > g/2$, and zero otherwise). Note that I use only half of gap g, since the merging driver takes into account the lead gap after the merge, and I assume that this is simply the half of the total gap (although it can be not exactly the case in reality).

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

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

Once I obtained the estimates for fixed $\bar{v}_F$, I can calculate ex-post late-merging penalty $C(v^B)$ according to 16. Note that critical gap $G_c(v^B, \bar{v}_F)$ is needed to calculate merging penalty; any appropriate model can be used for it, I use the model discussed earlier (see 13).

D Cellular-Automata Simulation Setup

In CA both time and space are discrete. In our simulation setup, I set the time step to 1 second, and I divide the lanes into cells of 5m each. Every cell can be either empty or occupied by a car, and every car has its own velocity, which is an integer number from 0 to v^{max} (although qualitatively similar, not to be confused with the v^{max} in our numerical experiment setup). The velocity indicates how many cells a car travels per time step. For consistency with our theoretical models, I allow lane-changing only for the cars on the blocked lane, and in order for a car to merge to the free lane, there must be enough space both ahead and behind.

Our CA-simulation setup is defined by the following rules:

i. Drivers on the blocked lane who want to merge, check if they have enough space available. For the merging to take place, the cell next to the car on the free lane must be empty. Moreover, there should be v empty cells ahead, where v is the velocity of the car, and v^{back} empty cells behind, where v^{back} is the velocity of the car right behind on the free lane;

ii. Every car on the blocked lane accelerates, or it is required to decelerate. That is, for every car I set $v \leftarrow \min\{v+1, x, v^{max}\}$, where x is a distance to the car ahead on the blocked lane;

iii. Every car with a positive velocity decreases the velocity by 1 with probability p. With a reasonable choice of p, this represents various stochastic real-traffic phenomena, such as over-braking or reaction delay at start.

iv. Every car moves ahead by v cells.

In our CA simulations I use an open-boundary condition: if the car reaches the end of the road segment while driving on the free lane, it exits the system. On the other hand, car arrivals are generated on the free lane in order to maintain a traffic flow with the constant density. When all blocked-lane drivers exit the bottleneck, the simulation terminates. While I could generate infinite inflow of B-drivers and terminating the simulation after some fixed time, I do not use this approach, since it prohibits us from meaningful use of total travel time metric. Indeed, if the number of B-lane drivers is not fixed, I would not know, whether a higher travel time is attributed to higher delays or simply due to a higher number of exited drivers (which is actually good). Hence, I randomly place a fixed number of B-lane drivers and record how much time they need to leave the bottleneck.

Once I have fixed all the parameters of the CA, I need to evaluate the parameters of the model in §3.1, so that I can obtain the optimal solution to the corresponding DP, and introduce it as a new rule in the CA framework. I could use parameter values that appear in the literature, or the ones derived by the numerical experiment setup in Appendix 2. Instead, I opt to use the CA simulation itself to evaluate these parameters, very similarly to what I would be doing if I were collecting real data. More specifically, the estimation method that I use is to repeat the CA simulations several times, and to take the sample averages of the quantities of interest as their estimates.

Velocity on the free lane, $\bar{v}_F$: unlike the model of §3.1, in a CA simulation, like in real life, the traffic flow on the free lane consists of many individual cars, each with its own position and velocity. For any given density of traffic flow in the free lane, I run repeatedly CA simulations (without any cars on the blocked lane), keeping track of the velocity of each car at all times. Averaging over the the time horizon, the different cars, and the various simulation runs, I obtain an estimate for $\bar{v}_F$;

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 §4 suggest that I choose $v_H = v^{max}$ and $v_L \approx \bar{v}_F$ (I set the low velocity, v_L, to the closest integer that is greater or equal than $\bar{v}_F$);

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

Late merging penalty P: viewed from a CA perspective, the late-merging penalty, P, in the model of §3.1 consists of different components. The first is the time spent during unsuccessful attempts to merge, while the car stands still near the blockage point. I denote this by P_0. Similarly to the estimation of q_H and q_L, I can estimate P_0 directly from CA simulations by averaging over many runs. I also introduce two analytically derived components. The final form of the late merging penalty is:

$$L = L_0 + C(0) + \frac{\bar{v}_F}{2} + \sqrt{2d},$$

if I express the distance in CA cells, and the velocity in CA cells per second. The second component is the merging penalty due to "relaxation". The third component is the travel time lost due to the required acceleration from $v = 0$ to $\bar{v}_F$ after merging. This detail is omitted in our theoretical model, but may play an important role when acceleration is limited. To see how I obtain it, note that according to the formula $x = (v_2^2 - v_1^2)/2a$, if I take $a = 1$ (since in CA this is the acceleration I use, expressed in cells per second squared), $v_2^2 = \bar{v}_F$ and $v_1^2 = 0$, I have the distance of acceleration $x = (\bar{v}_F)^2/2$. Trivially, time of acceleration is equal to $\bar{v}_F$ seconds. If the car would move with the constant velocity $\bar{v}_F$ for $\bar{v}_F$ seconds, it would travel $(\bar{v}_F)^2$ cells. The difference between these two distances is also $(\bar{v}_F)^2/2$, which is the distance lost due to the finite acceleration. This corresponds to $\bar{v}_F/2$ seconds of travel time lost due to the acceleration, since I assume that the flow on the free lane is moving with the constant velocity $\bar{v}_F$. The last component is the time required to cover a distance d after merging, and can be derived easily from the formula $at^2/2 = d$, if I again take the value for acceleration $a = 1$.

Merging penalties c_L and c_H: I use the theoretical model, described above in Appendix 2. As a distance function $d^F(\bar{v}_F)$ for free lane I use the one described by (14), and 13 as the model for critical gap $G_c(v^B, \bar{v}_F)$.

Apart from the aforementioned parameters, I 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 CA rules. In the latter, a driver can check for merging opportunities once per second. The distance between subsequent attempts clearly depends on the driver's velocity, so it is not constant. However, our theoretical model assumes that a driver checks for merging opportunities once every d meters. The best I can do, therefore, is to try to solve the single-car discrete-velocity DP for different choices of d. In §5 I only report on the case where $d = 20$m, since the results turned out to be quite robust to the choice of d.

Finally, I set the slowdown probability p for the CA to $p = 0.3$, which is a typical value in the literature, and I assume the length of the road segment to be equal to 500m.