

# **An Agent-Based Approach to Improving Commuter Traffic by Means of Advanced Traveler Information Systems**

---

A Thesis

Presented to

the Faculty of the School of Engineering and Applied Science

University of Virginia

---

In Partial Fulfillment

of the requirements for the Degree

Master of Science (Systems and Information Engineering)

by

Armen G. Melikian

December 2014

# Approval Sheet

This Thesis is submitted in partial fulfillment of the requirements for the degree of  
Master of Science (Systems and Information Engineering)



---

Armen G. Melikian

This Thesis has been read and approved by the Examining Committee:

**Matthew Gerber**

---

Advisor

**Donald E. Brown**

---

**William T. Scherer**

---

**Michael C. Smith**

---

**Gerard P. Learmonth, Sr**

Accepted for the School of Engineering and Applied Science:



---

James H. Aylor, Dean, School of Engineering and Applied Science

December 2014

# Abstract

Extensive urban and metropolitan development over the past decades has resulted in the need to better move goods and people over the nation’s motorways. However, physical limitations in motorway infrastructure and inefficient use of existing infrastructure have caused considerable economic loss through the late delivery of goods and loss of productive labor, as well as environmental damage, national security concerns, and financial burden resulting from fuel wasted during congestion.

Improving utilization of existing physical assets is a crucial and relatively inexpensive means to achieving those ends. The proliferation of inexpensive mobile devices over the past decade has fundamentally changed traffic management, both from the perspective of motorists, as well as planners and logisticians. Motorists now have rapid access to realtime traffic information, and can thus easily re-plan their journeys while already underway. Likewise, organizations that operate these services have access to motorists’ whereabouts and itineraries, and are in a position to shape and manage traffic by providing instructions to individual drivers.

This research addresses an exploratory agent-based approach towards modeling a road network in which agents 1. have varied routing behaviors and 2. may or may not have access to realtime rerouting capabilities. The objective is to explore what effect modifying driver portfolios will have on the “price of anarchy” associated with a commuter network. This information can be used by urban planners to determine which driver and information-access portfolios result in the least vehicle delay and best use of motorway resources, thus informing policies and incentives that can alleviate network congestion. Results on a Northern Virginia test case demonstrate that active rerouting has overall modestly positive effects on network latency for small and moderately-sized populations; however, in larger populations, active rerouting can lead to competitive pressures between agents, resulting in overall degradation of performance. It is suggested that future work implement competing managerial agent classes, as well as test over additional geospatial datasets.

# Contents

<b>Contents</b>	<b>4</b>
<b>1 Introduction</b>	<b>6</b>
<b>2 Problem and Objectives</b>	<b>10</b>
2.1 Problem Description . . . . .	10
2.2 Problem Definition . . . . .	11
2.3 Research Objectives . . . . .	12
<b>3 Literature Review</b>	<b>13</b>
3.1 Simulation Methods in Traffic Engineering . . . . .	13
3.1.1 Scope . . . . .	13
3.1.2 Methodological Approaches . . . . .	15
3.2 Infrastructure Capacity & Traffic Flows . . . . .	17
3.3 Game Theory and Traffic Networks . . . . .	18
<b>4 Methods</b>	<b>20</b>
4.1 Overview . . . . .	20
4.2 Platform and Software . . . . .	20
4.3 Geospatial Data . . . . .	21
4.4 Driver Classes and Behaviors . . . . .	23
4.5 Paramaters, Initialization, and Runtime . . . . .	25
4.6 Output and Metrics . . . . .	27
<b>5 Results and Analysis</b>	<b>28</b>
5.1 Experiment Design . . . . .	28
5.2 Results . . . . .	28
5.3 Model Calibration . . . . .	34
5.4 Verification and Validation . . . . .	36
<b>6 Conclusions and Considerations, Future Work, and Applications</b>	<b>39</b>
6.1 Conclusions and Contributions . . . . .	39
6.2 Considerations . . . . .	39
6.3 Future Work . . . . .	40
6.4 Applications . . . . .	41

<b>Bibliography</b>	<b>42</b>
<b>Appendix A</b>	<b>47</b>
<b>Appendix B</b>	<b>52</b>
<b>Appendix C</b>	<b>53</b>

# Introduction

Motorway congestion is a major cause of both material and immaterial costs for both commuters and commercial freight operations. Delays caused by road traffic are not only an inconvenience to drivers, but additionally are associated with economic loss, such as decreases in productive time resulting from protracted commuting, and the burning of fuel while idling or moving at reduced speeds. Commercial operations are also subject to the aforementioned costs; additionally, time-sensitive shipments depend on precise scheduling that can be disrupted by events that slow transit.

Texas A&M University's Transportation Institute compiles national transit and traffic data and releases its *Urban Mobility Report*<sup>1</sup> which details congestion statistics and trends, infrastructure changes and growth, and the environmental and economic effects of traffic patterns. Data is compiled for 498 urban areas in all 50 states and the District of Columbia, categorized as "very large" ( $x > 3$  million people), "large" ( $3 \text{ million} > x > 1 \text{ million}$  people), "medium" ( $1 \text{ million} > x > 500 \text{ thousand}$  people), and "small" ( $500 \text{ thousand people} > x$ ). The 2011 report states the following findings:

- The total economic cost of congestion of the 498 surveyed areas in 2011 was approximately \$121bn. This was the result of both wasted fuel and delay-related expenses (loss of productive hours, etc.). This amounts to 0.8% of the national GDP that year<sup>2</sup>.
- In 2011, Americans collectively spent 5.5bn hours stuck in traffic. This translates into 38 hours, on average, per individual. The national average of the Travel Time Index, which is the ratio of travel time during peak hours to travel time during free-flow conditions, was 1.18. This indicates an 18% increase in travel time during peak hours relative to free-flow travel time for a given commute (for example, a commute that

---

<sup>1</sup>Report available at <http://mobility.tamu.edu/ums/report/> [1]

<sup>2</sup>Via <http://data.worldbank.org/>

would be 30 minutes during free-flow conditions is instead over 35 minutes during peak hours).

- In 2011, Americans collectively wasted 2.9bn gallons of fuel as a result of road congestion. This resulted in an additional carbon dioxide output of 56bn pounds.
- The average monetary cost to an American commuter in 2011 was \$818, more than a threefold increase from 1980 (adjusting for inflation).
- Predictions indicate that in by 2020, the aggregate economic cost of congestion will increase to \$199bn (in 2011 dollars), with a total of 8.4bn hours wasted in traffic and 4.5bn gallons of fuel wasted. The cost to the individual commuter will be \$1,010 (again, in 2011 dollars) with 45 hours wasted in traffic.

Additionally, the report states that delays in truck deliveries - especially shipments required for “just in time” manufacturing - resulted in a cost of \$27bn (in 2011), much of which was passed onto consumers. This constitutes 22% of the above \$121bn value.

Moreover, unrelated research highlights the detrimental effects occupational commuting has on an individual’s mental and physical health; quality of sleep and chronic exhaustion were specifically linked to long, congested commutes [2]. While this has obvious implications with respect to quality of life, there also exist underlying public health concerns that could be considered preventable.

The *Urban Mobility Report* prescribes certain treatments; these range from expanding infrastructure to accommodate greater vehicle throughput, to encouraging changes in development patterns in a way that makes metropolitan areas less susceptible to jams. Of particular interest is a recommendation for developing management strategies that optimize transit given current infrastructure limitations.

## **Advanced Traveler Information Systems and the Intelligent Transportation Systems Strategic Research Plan**

Advanced traveler information systems (ATIS) are defined as any system that provide information intended to assist an overland vehicle operator in his or her planned journey. ATIS have been discussed in the transportation literature as far back as the 1990s; however, the rapid development and widespread proliferation of inexpensive and easy-to-use consumer mobile technologies in recent years have resulted in a substantial boon to transportation engineering, especially with respect to system-level infrastructure management. In 2009, the United States Department of Transportation published its Intelligent Transportation Systems (ITS) Strategic Research Plan 2010-2014<sup>3</sup>, a broad research initiative intended to improve roadway safety and mobility, as well as to ease the environmental impacts associated with road travel.

While the initiative covers many research areas over several disciplines, tapping consumer mobile technology and devices has been of considerable interest to the Department of Transportation, as well as the private sector. The ability to collect data in realtime from vehicle operators, as well as broadcast realtime information to said operators has the potential to aid logisticians in managing vehicle traffic and making optimal use of road infrastructure. The ability to both access and broadcast data of such high resolution (down to the individual vehicle) can allow planners and traffic engineers to improve the regular use of roadways, as well as efficiently and appropriately respond to highly disruptive events and conditions.

## **Research Summary**

An agent-based simulation has been employed to run a number of experiments over a simulated “rush hour,” approximately three real-time hours<sup>4</sup>, in which parameters pertaining to total vehicle population, access to ATIS, and routing preferences have been varied. The problem is formally defined in Section 2. How experiments have been designed and run are explained in Sections 4 and 5. From these experiments, results pertaining to mean vehicle

---

<sup>3</sup>[http://www.its.dot.gov/strategic\\_plan2010\\_2014/index.htm](http://www.its.dot.gov/strategic_plan2010_2014/index.htm)

<sup>4</sup>According to Texas A&M’s *Urban Mobility Report*, a typical morning in a rush-hour in a large/very large metro area lasts approximately three hours.



travel time (compared to a hypothetical, unconstrained free-flow case) are analyzed and implications are interpreted (Sections 5 and 6).

# Problem and Objectives

## 2.1 Problem Description

On any given day, motorists will commute to and from their places of employment (or other destinations); the time it takes to travel from one's origin to destination is dependent on the vehicle operator's driving behavior and route selection choices, the driving behavior and route selection choices of other drivers, the number of drivers on the road at any given time, traffic accidents, non-vehicle obstructions, inclement weather, road closures, and acts of God. In effect, travel outcomes are determined by one's own actions, the actions of others, and externalities unrelated to agent decision-making.

Given the pervasiveness of smartphones and similar technologies, it is now possible for services such as Google Maps, Waze, Nokia HERE, and INRIX to track, in realtime, where, when, and under what circumstances delays and congestion occur with high spatiotemporal precision. Moreover, the information gathered by these services can be disseminated back to their users, who can revise their routes to shorten travel time. This continuous realtime feedback process can be used to better manage aggregate motorway traffic and shorten travel time for individual vehicles.

However, executive functions are still left to the vehicle operator, so actual moving behavior may occur irrespective of what information is available to any individual driver; under real conditions, that decision is typically opaque. In other words, a driver may behave in a way that seems contradictory to the routing information provided to him for reasons unknown to a third-party or outside observer.

The work documented here explores several test cases in which drivers on a network exhibit different moving behaviors/preferences, and may or may not heed to route alternatives and/or revisions. Most importantly, their interactions with one another do indeed affect their

own transit outcomes. For each test case, performance against a hypothetical best case is compared - this is further explained below.

## 2.2 Problem Definition

This research aims to simulate commuter traffic over a road network by means of a mesoscopic traffic simulation built as an agent-based model. This simulation is designed to demonstrate the “price of anarchy” on a road network - that is, how self-interested drivers with different itineraries and travel strategies (utility preferences) independently interact with one another and with the network itself, given varied levels of awareness of the network’s state. More formally, this is defined as:

1. Agents are self-interested and seek to minimize their travel time (latency)<sup>1</sup>.
2. Agents will select a path based on factors that are not time-variant.
3. Agents are non-cooperative and have imperfect information on the state of the system.
4. A small number of agents have perfect information on the state of the system at all times, and are willing (to a degree) to adjust their planned routes to avoid high-use areas at the expense of distance traveled.
5. Given these conditions, determine how the partial population of the agents described in (4) impacts mean latency. Additionally, determine whether these outcomes are bounded or amplified by scale.

Observations pertaining to vehicle travel time and network congestion, given simulation parameter settings, can be used by urban planners and entities providing navigation services to inform policy and determine appropriate incentive strategies.

---

<sup>1</sup>“Latency” as defined here is equivalent to the US Department of Transportation Federal Highway Administration’s term *planning time index* [3]. “Latency” is the term used in the game theory literature, and will be used throughout this document

## 2.3 Research Objectives

The objective of the documented research is to investigate how different driver portfolios impact network congestion in a test case representative of a real-world motorway network. More specifically, the goal is to determine how varying the number of drivers with access to perfect information of the state of the system (as a fraction of the total population) affects travel latency for both types of drivers. Of interest are whether effects are independent of total test population sizes, whether one observes diminishing or increasing returns as population fractions are adjusted, and if there exist any cases in which one observes a degradation in performance in absolute terms.

The work is exploratory in nature and it is intended that methods and results presented here can be used to inform high-level policy for urban planners, logisticians, and private services such as Google Maps, Waze, Nokia HERE, and INRIX. This is especially useful in the event that major changes are planned to physical infrastructure, in that systemic effects can be simulated beforehand.

More powerful (and perhaps less obvious) is the potential for providing actionable guidance to urban planners to the effect of informing policy as it pertains to driver incentives and future development. Coupled with detailed information on commuter populations (obtained through surveys, for example), urban planners can better engineer incentives and infrastructure development to suit local communities. Both of the above are more thoroughly explained in Section 6.

# Literature Review

The majority of the literature review pertained to traffic simulation, addressing scope, methodology, and their respective case-by-case utility. Supplemental research was done on traffic flow theory and game theory to better inform simulation mechanics. The research put forth in this document draws from all three, but intends to contribute primarily to the first.

## 3.1 Simulation Methods in Traffic Engineering

### 3.1.1 Scope

#### Microscopic Traffic Simulation

A microscopic traffic simulation (or microsimulation) is one in which individual vehicles are modeled and observations are made at the unit-level. Microsimulations parameterize individual vehicles' position, velocity, and acceleration, and are often used to simulate car-following, acceleration/deceleration behavior, and lane-changing behaviors. They are useful in modeling how road features such as stoplights, bottlenecks (lanes merging), and dedicated turn lanes can affect vehicle movement and local congestion/traffic patterns. Due to large computational overhead, it is typically impractical to simulate a large area using a microsimulation [4–7]. Several microsimulation platforms exist, notably TRANSIMS, a project originally developed by Los Alamos National Laboratory; it is unique in that it can simulate populations using census demographic data, adding a level of sophistication not present in other platforms [8].

## Macroscopic Traffic Simulation

A macroscopic traffic simulation (or macrosimulation) is one in which aggregate system dynamics are modeled, as opposed to individual units. Such simulations explicitly **do not** model individual vehicles, their movement, or their interactions with one another. Instead, they model and output aggregate metrics such as average vehicle speed, traffic flow, and traffic density over a surveyed area. Such simulations are typically built on fluid dynamic models. Because of their low resolution, they are limited to use on larger road networks or highway systems, where unit-level detail is not necessarily required. However, even in this application they are limited, since driver behaviors or distinct driver types that can affect system-wide outcomes are not distinguished from one another [5, 9, 10].

## Mesoscopic Traffic Simulation

Very generally, mesoscopic traffic simulations fall between macroscopic simulations and microscopic simulations in terms of resolution and actionable scope [5, 9]. Mesoscopic models will typically model individual units (like microsimulations), or instead small groups of units (“platoons”), providing higher resolution than would be available in a macrosimulation. However, high-resolution vehicle *behaviors* characteristic of microsimulations, such as car-following and lane-changing, are typically absent/omitted, or at least vastly simplified. Mesoscopic models can be applied to larger areas/populations than would be practical for a microsimulation, while affording a level of detail absent from a macrosimulation.

Much of the traffic engineering literature on mesoscopic simulation has involved building hybridized mesoscopic/microscopic models [11–13]. These models are intended to model and generally characterize large networks as accurately as possible, while allowing for high-resolution modeling for specific segments of the network that are considered to be of greater interest. Recent literature has also seen this approach applied to logistics analysis, in that it is computationally less expensive than microsimulation, and is adequately detailed for larger scale requirements and longer travel itineraries [14].

### 3.1.2 Methodological Approaches

#### Fluid Dynamics

Fluid dynamic and gas-kinetic models (derived from the physical sciences) have been applied to macroscopic traffic simulation since at the 1950s. Lighthill and his former student, Whitham [15] applied a kinematic wave model to describe traffic flow and bottlenecking on single lane “crowded arterial roads.” Richards [16] arrived at this same conclusions as Lighthill and Whitman and published his findings in 1956. The characteristic equation of both publications includes the Greenshields expression, which was formulated in the 1930s to relate vehicle speed and travel density (see Section 3.2) [9]. Subsequently, Whitham and Payne [17] [18] added a second equation to the 1955 publication, and the two together named for Lighthill, Whitman, Richards, and Payne have since then been featured (and solved for) in the traffic engineering literature. This approach employs a system of partial differential equations, for which exact solutions exist [19].

#### Discrete Event Models and Cellular Automata

The Nagel-Schreckenberg model [20] is a well-known cellular automaton that models traffic flow. A single-lane road is divided into segments (cells) which have a state of either occupied or unoccupied. Vehicles execute a simple set of instructions: a car not at the maximum speed allowed by the model will accelerate, it will then reduce speed based on their current speed and the distance between it and the vehicle ahead of it, it will then reduce its speed with some random probability  $p$ , and finally it will execute its movement forward. Travel is in one direction only, and vehicles cannot overtake one another. This model was noteworthy because proved traffic flow was an emergent phenomena, and modeled it at a very high granularity. The aforementioned TRANSIMS model was developed by Nagel at Lawrence-Livermore, and is based largely on the Nagel-Schreckenberg model [8]. The Nagel-Schreckenberg model has been further extended to more complex, multi-lane simulations [21, 22].

Discrete event simulation has been independently employed as a way to model traffic flows and vehicle movements. Burghout et. al. [11, 13] employ a discrete-event methodology in the aforementioned hybrid micro/mesoscopic traffic simulation, in which vehicle movement is modeled as an arrival/queuing process.

## **Agent-Based Models**

Agent-based models have been widely employed in the field of traffic engineering to investigate a wide variety of conditions and involved factors [23–25]. Human choice and decision-making behaviors, route guidance, traffic flow dynamics, and broad systemic phenomena have all been approached and analyzed using ABMs at various levels of complexity and scope.

Most relevant to the research here are works by Hussein Dia [26], Wahle et. al. [27], and Wahle et. al.[28], who explore the impact of real-time traffic information provided to drivers (such as by means of as an ATIS) using ABM and cellular automata. Whale et. al. are concerned more with the systemic effects associated with ATIS utilization, specifically what time of information is made available to motorist, and agent behavioral coding is fairly simple. Dia employs a rather sophisticated agent utility framework derived from demographic data and surveys distributed to human motorists in Brisbane, Queensland, Australia. This provides a highly illustrative simulation environment, with respect to both agents and the infrastructure on which they operate. The research proposed in this document is in a similar vein to the three papers described above, with focus on the health of the overall infrastructure network. Work by Buscema et. al. [29] features a simulation built in NetLogo that models Braess’s Paradox with ATIS augmentation; this is perhaps the most obvious simulation publication surveyed with a direct analog to network equilibrium in game theory (Section 3.3). Balmer et. al. demonstrate that ABM is well-suited for mesoscopic simulations [30].

Of methodological importance is also the platform used for simulation purposes. Multiple ABM and platforms have been employed in traffic engineering for simulation purposes, each with its own strengths and weaknesses pertaining to model complexity and size, speed, and



ease of use. Of note are NetLogo<sup>1</sup> (a Logo derivative developed at Northwestern University) and Repast Simphony<sup>2</sup> (developed at the University of Chicago). In addition to ABM and CA platforms, there exist several traffic-specific simulation toolkits, both closed-source proprietary and open-source.

## 3.2 Infrastructure Capacity & Traffic Flows

Motorway capacity and traffic flow has typically been described using the fundamental diagram of traffic flow. Development began with work by civil engineer Bruce Greenshields, first appearing in a publication in 1933 [31, 32]. In the decades since then, the fundamental diagram has continued to serve as the basis on which traffic flows and motorway congestion are modeled, with some modifications and statistical validation in various case studies. The diagram is typically expressed in the form of:

- **Flux-Density:** The relationship between flux (number of cars per hour) and density (number of cars per mile).

A maximum exists in the relationship between flux and density, separating a stable free-flow regime from an unstable congested flow regime. This maximum is considered the roadway capacity. The diagram (or rather, the relationships represented in it) can also be expressed as:

- **Speed-Density:** The relationship between speed (miles per hour) and density (number of cars per mile).
- **Speed-Flux:** The relationship between speed (miles per hour) and flux (number of cars per hour).

These alternate forms are often found and used in the traffic engineering literature. Figure 3.1 below illustrates all three forms.

---

<sup>1</sup><http://ccl.northwestern.edu/netlogo/>

<sup>2</sup><http://repast.sourceforge.net/>




Figure 3.1: The fundamental diagram of traffic flow. From upper-left, clockwise: speed-density, speed-flux, flux-density. Critical points are demarcated by faded lines indicating changes in flow regime. Of most interest is the figure in the bottom left. To the left of the critical point one experiences free or bounded flow. To the right, one experiences jams. Capacity is measured by flow at the maximum [30].

While not directly included in the research completed here, the fundamental theory of traffic flow (from which the fundamental diagram is derived) is worth mentioning due to its legacy and importance in traffic engineering, especially as it pertains to infrastructure capacity.

### 3.3 Game Theory and Traffic Networks

Traffic networks and routing have been examined in a game theoretical paradigm for several decades, providing a rigorous theoretical basis on which to build a simulation (as is being done here). In a paper published in 1952, John Glen Wardrop put forth two equilibrium principles stating that self-interested drivers will non-cooperatively seek to minimize their travel costs

and that under equilibrium no driver can by himself lower his own cost, and that the average journey time of all drivers is minimized at equilibrium [33]. Beckmann demonstrated that traffic flows on a network are satisfied by a unique Nash equilibrium [34]. Tsitslikis and Bertsekas demonstrate that such equilibriums would arise over a distributed shortest-path protocol [35]. Koutsoupias and Papadimitriou coin the term “price of anarchy” to describe the detrimental effects of noncooperative behavior [36]. Both they and Roughgarden and Tardos demonstrate that selfish routing choices by agents on a transportation network have a quantifiable impact on network performance [37]. One historically well-known example of this is Braess’s paradox, in which cost-minimizing agents will end up *reducing* overall network performance with the *increase* of network capacity [38, 39] by means of selfish routing behavior. Roughgarden and Tardos [37], Koutsoupias and Papadimitriou [36], and Papadimitriou and Valiant [40] independently present solutions to the performance degradation resulting from selfish behavior. Roughgarden has independently investigated network problems associated with selfish choice-selection in game theory [41]. Additional work has been done in applying the price of anarchy to mixed games (“mixed price of anarchy”) [42, 43] and games of imperfect information (“Bayes-Nash price of anarchy”) [44, 45].

# Methods

## 4.1 Overview

As described in Section 2, this research investigates how drivers with different preferences, with access to varied amounts of information, affect the system-wide performance of a traffic network. This is done by means of a mesoscopic traffic simulation (Section 3.1.1), built using an agent-based modeling platform. This approach was selected because, unlike others, it is *agent-oriented* as opposed to *process-oriented*; the interactions of individual entities with one another, and with their environment, allows for emergent outcomes that would not necessarily occur with a process-oriented or deterministic approach.

## 4.2 Platform and Software

The simulation developed for this project was built using NetLogo 5.1.0<sup>1</sup>. NetLogo is a dedicated agent-based modeling language and IDE written in Java and Scala. It is entirely open-source and free to use, and is popular in academia as both a pedagogical resource and research tool. NetLogo supports a number of extensions that allow for interfacing with other programming languages and multiple data formats. Simulations written in NetLogo have appeared in the traffic engineering literature, and NetLogo comes bundled with a number of simple traffic simulation models.

Geospatial data was processed using ArcMaps 10.1 of the ArcGIS suite, produced and maintained by ESRI.

---

<sup>1</sup>The installation used to write and run the simulation was modified in order to increase the memory allocation pool in the Java VM. Attempting to run this simulation without increasing the memory heap will result in a runtime error. Instructions available here: <http://netlogo-users.18673.x6.nabble.com/Pointing-NetLogo-to-64-bit-JVM-td5002961.html>



modeling area. The map package was converted into .shp format, and two modifications were made to the output to reduce computational complexity. They are as follows:

- Removal of urban routes, secondary routes, and frontage routes. The simulation environment includes only interstate highways, VA state highways, US highways, and their respective on/off-ramps. This was done to reduce computational complexity in the simulation environment. Likewise, it also reflects that the majority of a high-mileage commute will occur on highway (as opposed to residential) roads.
- Consolidation of distinct travel directions. Throughways with distinct travel directions that are separated by a barrier or median (such as most highways) are illustrated as such.




Figure 4.2: A segment of I-66 from LRS13.4. Notice that travel directions are illustrated separately.

Southbound and westbound features were removed from the shapefile. This was done to both reduce computational complexity, and to account for limitations in discerning between directional features in NetLogo (explained below).

The modified shapefile was then imported into NetLogo using NetLogo’s GIS extension. The environment was built using a “node-and-link” approach; the ends of polyline features are represented by a breed of non-moving turtles (called “nodes”), and the polylines themselves are represented by links between said nodes.

The absence of southbound and westbound travel directions was corrected for by representing each polyline feature as a pair of opposite-facing directed links. This simplification assumes structural symmetry between opposite travel directions. This approach was selected instead of single undirected links for the purpose of tracking which and how many vehicles were traveling in each direction.




Figure 4.3: A segment of highway from LRS13.4 imported into NetLogo. Nodes represent polyline ends, and each directed link represents one travel direction. The nodes are color-coded to reflect road type consistent with the LRS coding. Blue indicates Interstate highways, red indicates US highways, and black indicates state highways.

## 4.4 Driver Classes and Behaviors

Drivers are divided into three classes that reflect different route selection behaviors. They are as follows.

- **Stubborn** Stubborn vehicles take the shortest point-to-point path from their origin to destination. A stubborn vehicle’s route cannot be altered once it is initialized. This shortest path has the option to be weighted based on road type, in which case a coefficient selected by the simulation operator can be used to penalize (and therefore discourage) links with a multiplicative coefficient. For example, the user could choose to not penalize interstate features (coefficient of one), penalize US highways with a coefficient of 1.5 (100 meters of US highway is of equal “score” to 150 meters of interstate), and penalize VA state highways with a coefficient of 2. This would be done to reflect capacity and throughput asymmetry, resulting from features such as the number of lanes, the presence of traffic lights and intersections, etc. By default, all coefficients are set to 1, indicating no bias or preference by the simulation operator.
- **Stupid** Stupid drivers take a “random” path. A link’s score is the product of its length and a uniformly-distributed random variable on the interval [1, 10]. Once a stupid vehicle is initialized its route cannot be changed. While these vehicles may not realistically be observed in the real world, their purposes is to further introduce a degree of randomness into a system that is fundamentally anarchic.
- **Smart** Smart vehicles are functionally ATIS-equipped (and compliant). Smart vehicles have complete information regarding the state of the traffic network (including the whereabouts of both stubborn and stupid vehicles), and will take a weighted shortest path that sacrifices travel distance to avoid high-traffic areas. The score of any given link is computed as follows:

$$s_L = d_L + ck_L$$

Where  $s_L$  is the link’s score,  $d_L$  is the length of the link,  $k_L$  is vehicle the density of the link, and  $c$  is an empirically-derived constant that, given high density traffic conditions, approximately doubles the score of a link with  $d_L$  equal to the mean link length of



the entire network (in other words, given heavy vehicle density,  $ck_L \approx \hat{d}$ ). This is recomputed at each tick, and vehicles of this class are rerouted accordingly; that is, a vehicle’s path is updated while it is already underway.

Paths are found using NetLogo’s bundled NW-Extension<sup>2</sup>, which employs a modified version of Dijkstra’s algorithm.

Speed is set and adjusted based on the local density of its current link. Calibration is explained in Section 5.3.

## 4.5 Paramaters, Initialization, and Runtime

The simulation environment is initialized by importing the GIS shapefile and corresponding projection file, which is converted into a list of coordinate pairs, and subsequently built using node-and-link construction in the NetLogo world<sup>3</sup>. Road type classifiers (interstate, US highway, VA state highway) are preserved during this stage and features are assigned accordingly. Additionally, penalty coefficients for each road type can be assigned by the user at this stage.

Subsequently, a user-specified number of each vehicle type is initialized. These first-generation vehicles are assigned a random origin on the network, a destination (this is always a leaf node, i.e. a node connected to only one other node), and a path between the two dictated by the vehicle’s type.

NetLogo employs tick-based updating, where a tick is an arbitrary, discrete unit of time. Vehicles follow their assigned path<sup>4</sup>, adjusting speed as a function of the current link’s vehicle density during that tick. Every vehicle’s speed and every link’s vehicle density are recorded and stored as agent attributes during each tick. Vehicles are “killed” upon reaching their destination. At the beginning of every tick following initialization, each vehicle population is

---

<sup>2</sup>NW-Extension was first bundled with version 5.1.0 and replaces the previous Network-Extension.

<sup>3</sup>Code from Newcastle Cycling Experiment [46] used with permission.

<sup>4</sup>Code from Axwoman [47] used with permission, as well as Link Walking Turtles [48] and Move Towards Target [49] from the bundled NetLogo models library.

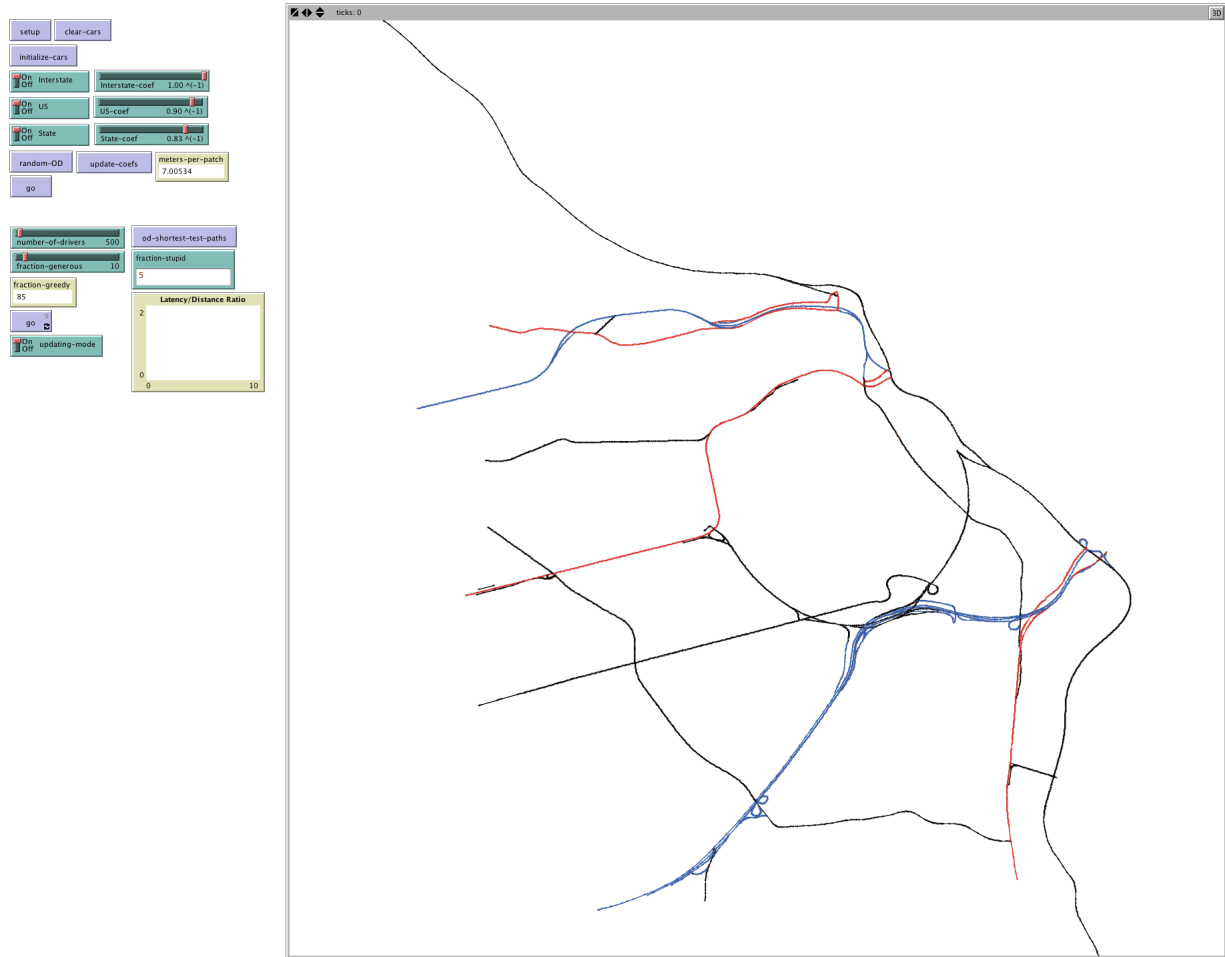


Figure 4.4: Netlogo Interface featuring all parameters and the simulation environment.

counted, and new vehicles are initialized at leaf nodes to keep the overall population and constituent subpopulations constant. As stated previously, ticks are an arbitrary unit of time that do not necessarily translate to a real-world value; the model is calibrated such that, given unrestrained traffic flow, a vehicle will travel at what would be 65 miles per hour. The model is updated and agent attributes recalculated what would be once every thirty seconds in realtime.

## 4.6 Output and Metrics

At the end of each tick, mean, maximum, and minimum speed of each vehicle population (and overall aggregate) is reported and plotted. In addition, the simulation reports and records mean vehicle latency. Game theorists writing on the “price of anarchy” measure the cost of selfish routing on networks as the ratio of the Nash equilibrium latency to the best-case, non-congested latency [37, 38, 41, 43, 45]. The latency ratio employed as a metric here is similarly defined to the aforementioned “price of anarchy”; it is the ratio of the in-simulation observed travel time to the best-case shortest-path origin-to-destination travel time given a bounded free-flow speed (practically speaking, a speed limit). Under ideal circumstances (or if the vehicle is the only one present), this value approaches one. In practice, it is greater than one.

The latency ratio is a useful metric for stubborn vehicles, since their routes are locked-in when a vehicle is initialized and the only variable during the vehicle’s commute is its speed. However, this is not the case with smart vehicles, as both their speeds and routes will change as they travel. To compensate this, a third metric is introduced specific to smart vehicles: a ratio of the vehicle’s actual path length over the length of shortest path between the origin and destination. Under ideal conditions this would be equal to one. In practice it tends to be greater than one. This metric can be compared against the latency ratio to quantify how much a smart driver has been “inconvenienced” from a milage standpoint to avoid high-density areas.

# Results and Analysis

## 5.1 Experiment Design

A total of 2475 experiments were run on populations of 500, 750, 1000, 1500, 2250, 3000, 5000, 7000, and 10000 total vehicles (225 experiments per population). The simulated trial period was equivalent to three realtime hours. In all cases, the stupid vehicle population was set to 5% of the total to introduce a low level randomness or “noise” into the system. The smart vehicle population was varied between 0% and 25% in increments of 2.5%. This amounted to twenty-five trials per permutation. Penalty coefficients were set to 1 for interstate highways, 1.1 for US highways, and 1.2 for state highways in all cases. These values were determined through heuristic trials to discourage use of lower-capacity options (US and VA state highways) without causing unrealistically long detours. Maximum free-flow speed was set to 110 NetLogo world-coordinate units per tick, or what would be approximately 65 miles per hour in reality.

## 5.2 Results

Analysis of mean vehicle speed indicates a positive correlation between fraction of smart drivers present and an increase of mean vehicle speed over the entire system. The impact of smart drivers is increasingly pronounced in larger total populations, with a non-trivial performance roll-off in the largest populations. Numeric results are presented in Figure 5.1, as well as the constituent graphs in Appendix A.

In all cases, the most substantial relative increase in overall mean vehicle speed occurs either between zero and 2.5% or 2.5% and 5%. Improvements still accrue as the fraction of smart drivers is increased, however there appears to be a roll-off in efficacy as the fraction

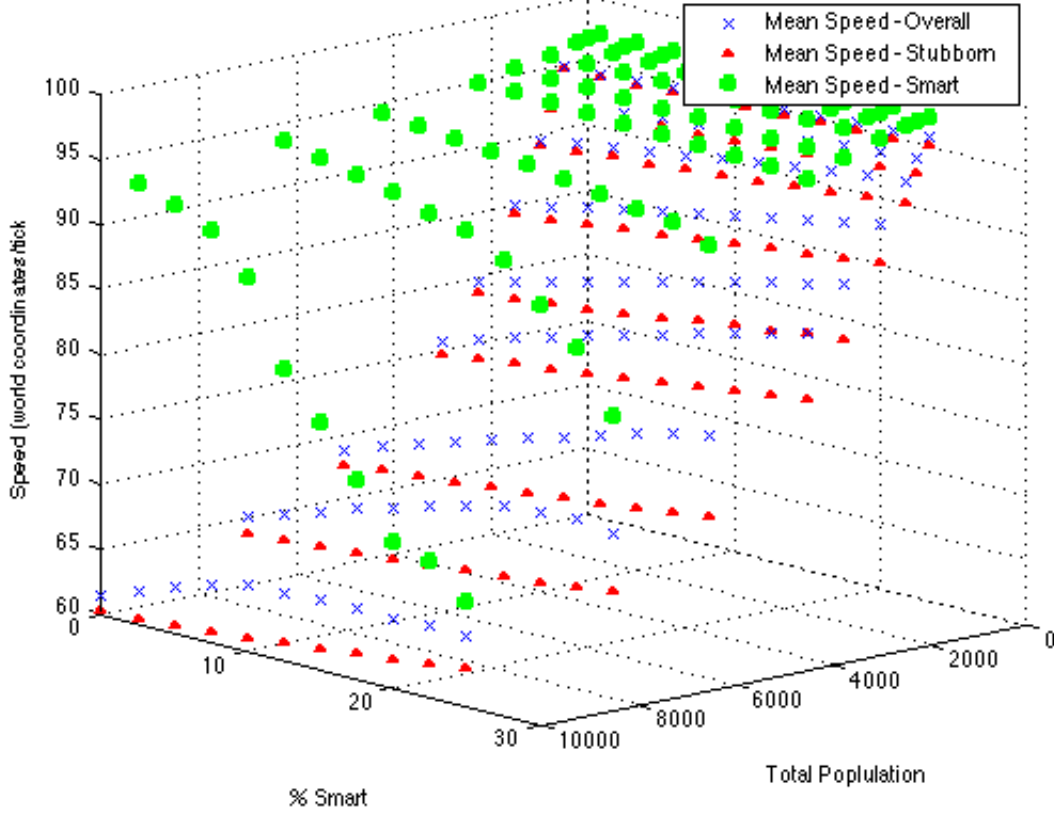


Figure 5.1: Mean vehicle speed plotted as a function of total population size and partial population of smart drivers. Individual population plots for the sake of clarity are presented in Appendix A.

is increased. Each 2.5% increase of smart drivers resulted in, on average, an increase of mean speed of 0.181% for the 500 driver experiments, 0.323% for the 750 driver experiments, 0.433% for the 1000 driver experiments, 0.643% for the 1500 driver experiments, 0.841% for the 2250 driver experiments, 0.984% for the 3000 driver experiments, 1.13% for the 5000 driver experiments, 0.875% for the 7000 driver experiments, and 0.611% for the 10000 driver experiments. Between zero and 25% smart driver fraction, this amounts to a 1.82%, 3.2%, 4.42%, 6.62%, 8.73% 10.3%, 11.9%, 9.08%, and 6.26% increase in speed respectively.

A few trends should be noted. First, the mean speed of stubborn vehicles monotonically increases with each fractional increase of smart vehicles. In no instance does the mean speed of stubborn vehicles decrease in any population size, even as the overall mean speed

% Smart	500	750	1000	1500	2250	3000	5000	7000	10000
0%-2.5%	0.1951	0.4472	0.4552	0.7298	1.0111	1.1278	1.358	1.4613	1.5479
2.5%-5%	0.2263	0.3647	0.5551	0.7432	0.8732	1.0964	1.3172	1.3463	1.564
5%-7.5%	0.2042	0.3455	0.4188	0.6515	0.8428	1.0689	1.3158	1.3877	1.3703
7.5%-10%	0.1925	0.3124	0.4308	0.6048	0.9095	1.0609	1.1975	1.225	1.1116
10%-12.5%	0.2242	0.2743	0.481	0.6791	0.8559	0.9891	1.1724	1.1523	0.1748
12.5%-15%	0.1242	0.3246	0.4068	0.6613	0.9023	0.8984	1.0572	1.1811	0.4173
15%-17.5%	0.1857	0.3048	0.4495	0.6109	0.8081	0.952	1.1288	0.8145	0.0677
17.5%-20%	0.1677	0.3126	0.4257	0.6079	0.7400	0.9065	1.0473	0.4241	-0.3335
20%-22.5%	0.1395	0.2513	0.3883	0.5609	0.7922	0.8624	0.9623	0.3009	0.3586
22.5%-25%	0.1464	0.2895	0.32	0.5778	0.6704	0.8776	0.7804	-0.5433	-0.1657

Table 5.1: Percentage change in overall mean speed at 2.5% increments of smart partial population. Column headers indicate overall vehicle population (500 cars, 750 cars, etc.). Tables for stubborn and smart vehicles alone are available in Appendix B.

or smart mean speed levels off or worsens (as seen in 7000 and 10000 vehicle populations. Second, there exists a performance roll-off as a function of population size when increasing the fraction of smart drivers. As population increases, the relative improvement experienced by each 2.5% increase in the smart vehicle population increases up to populations of 5000 vehicles. Beyond 5000 vehicles, that relative improvement begins to worsen. Third, the smart vehicle mean speed remains approximately stagnant regardless of their proportion to the total population, up to overall populations of 5000. In 5000 vehicle populations, smart vehicles travel approximately 4% slower when representing 25% of all cars than they do when representing 2.5%. In 7000 vehicle populations, they travel 17% slower over the same range. In 10000 vehicle populations, that change is even worse, at 27%. In fact, smart vehicles in this population travel only 8% faster than stubborn vehicles.

Latency results indicate that populations with no smart drivers do no better than a latency ratio of 1.65 (trips are roughly 65% longer time-wise than they would be at free-flow conditions). Introduction of smart drivers improves the latency of remaining stubborn drivers by 0.170%, 0.263%, 0.346%, 0.408%, 0.459%, 0.456%, 0.399%, 0.314%, and 0.321% for vehicle populations of 500, 750, 1000, 1500, 2250, 3000, 5000, 7000, and 10000 for every 2.5% fractional increase. Overall latency (across all driver types) was improved by 0.200%, 0.313%,

0.420%, 0.526%, 0.615%, 0.626%, and 0.498% per 2.5% increase, respectively, for populations up to 5000; for 7000 and 10000 vehicle populations, overall latency actually worsened by 0.179% and 0.945%, respectively. In these two populations, modest improvements averaging 0.344% and 0.341% are observed for smart vehicle fractions between 2.5% and 17.5%, and 2.5% and 7.5%, respectively; beyond those proportions, overall values experience degradation. Between zero and 25% smart driver fraction, greedy vehicles enjoyed a 1.69%, 2.60%, 3.40%, 4.00%, 4.49%, 4.47%, 3.93%, 3.10%, and 3.17% improvement in vehicle latency for the previously stated range, and the improvement across all vehicles was 1.98%, 3.09%, 4.12%, 5.14%, 5.98%, 6.09%, and 4.87% up to and including 5000 vehicles, with overall degradation to the effect of 1.76% and 9.81% in 7000 and 10000 vehicle populations.

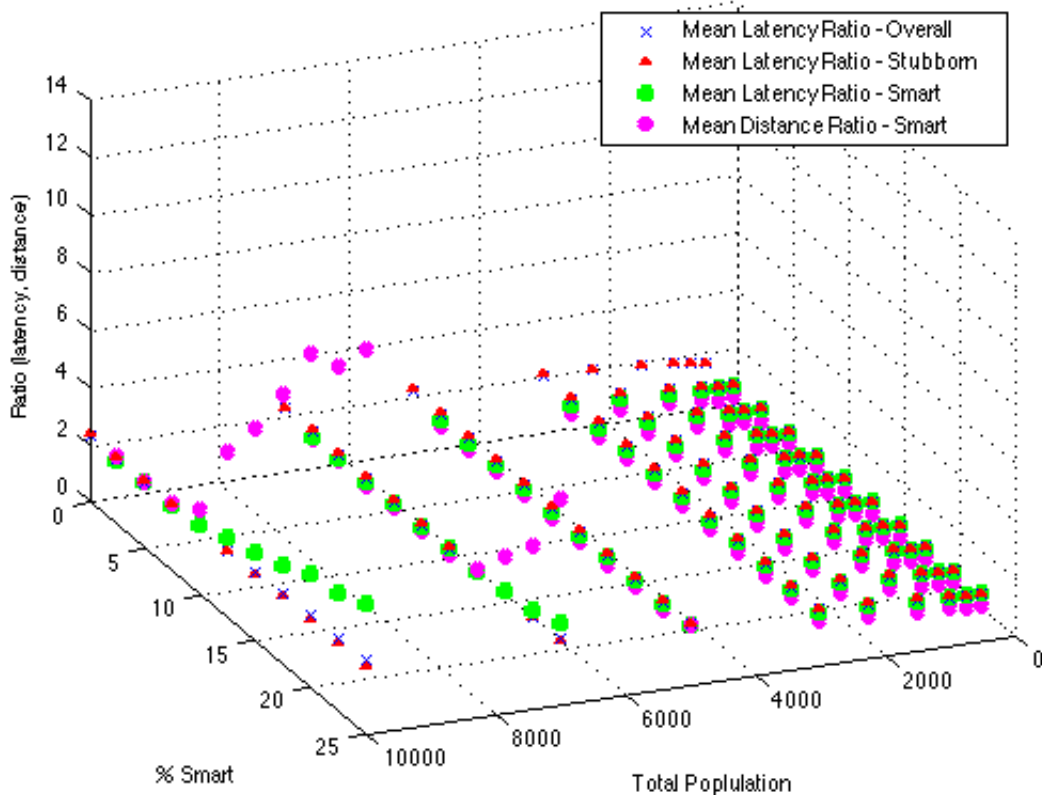


Figure 5.2: Mean vehicle latency plotted as a function of total population size and partial population of smart drivers. Individual population plots for the sake of clarity are presented in Appendix C.

% Smart	500	750	1000	1500	2250	3000	5000	7000	10000
0%-2.5%	-0.1475	-0.4344	-0.5474	-0.6644	-0.8512	-0.8107	-0.66	-0.5035	-0.4162
2.5%-5%	-0.3587	-0.3852	-0.4476	-0.5904	-0.6477	-0.6803	-0.6248	-0.4928	-0.343
5%-7.5%	-0.3312	-0.3936	-0.5177	-0.5227	-0.5597	-0.6501	-0.5772	-0.375	-0.265
7.5%-10%	-0.0426	-0.2151	-0.4394	-0.5875	-0.6459	-0.736	-0.5439	-0.3731	0.127
10%-12.5%	-0.1935	-0.3112	-0.4406	-0.5238	-0.6891	-0.6127	-0.5596	-0.2889	1.3543
12.5%-15%	-0.287	-0.3536	-0.4156	-0.5257	-0.5891	-0.5298	-0.4445	-0.3391	1.134
15%-17.5%	-0.0974	-0.3157	-0.3364	-0.5239	-0.6321	-0.6115	-0.5018	-0.0362	1.7896
17.5%-20%	-0.2234	-0.2012	-0.4306	-0.4946	-0.4741	-0.5735	-0.516	0.787	2.4458
20%-22.5%	-0.0897	-0.3151	-0.2455	-0.3852	-0.5802	-0.5401	-0.4158	0.9339	1.3188
22.5%-25%	-0.2229	-0.203	-0.3798	-0.4407	-0.4825	-0.5185	-0.1378	2.4741	2.3062

Table 5.2: Percentage change in mean overall vehicle latency at 2.5% increments of smart partial population.

% Smart	500	750	1000	1500	2250	3000	5000	7000	10000
0%-2.5%	-0.0562	-0.2834	-0.3706	-0.4086	-0.5379	-0.4906	-0.3184	-0.242	-0.2874
2.5%-5%	-0.3097	-0.2744	-0.292	-0.381	-0.3752	-0.3762	-0.3441	-0.2939	-0.2744
5%-7.5%	-0.2423	-0.3461	-0.3967	-0.3441	-0.3298	-0.4008	-0.3557	-0.2376	-0.2819
7.5%-10%	-0.0461	-0.1178	-0.3461	-0.4319	-0.4062	-0.4949	-0.3357	-0.2783	-0.3148
10%-12.5%	-0.1609	-0.2377	-0.342	-0.3604	-0.5482	-0.4357	-0.4069	-0.3163	-0.2774
12.5%-15%	-0.2699	-0.3519	-0.3486	-0.4427	-0.4456	-0.3413	-0.3563	-0.3323	-0.3664
15%-17.5%	-0.008	-0.2644	-0.3261	-0.4627	-0.5246	-0.5163	-0.4214	-0.342	-0.3425
17.5%-20%	-0.293	-0.2305	-0.4241	-0.4262	-0.414	-0.5036	-0.5054	-0.3557	-0.2994
20%-22.5%	-0.0503	-0.3134	-0.2245	-0.3628	-0.5312	-0.4996	-0.4313	-0.3429	-0.3645
22.5%-25%	-0.2626	-0.2144	-0.3866	-0.4548	-0.4737	-0.505	-0.5245	-0.4066	-0.4082

Table 5.3: Percentage change in mean stubborn vehicle latency at 2.5% increments of smart partial population.

As stated previously, latency is more difficult to interpret (and not particularly useful on its own) for smart vehicles, as they are constantly being rerouted while already underway to avoid congested areas. Thus, they will likely incur a distance penalty that can be substantial. To account for this, another metric, the distance ratio, was recorded for these vehicles - the ratio of actual traveled distance to the point-to-point distance had they instead been stubborn. Mean distance ratio was 1.18 for a total population of 500 (meaning travelled routes were roughly 18% longer than the shortest path), 1.22 for a total population of 750, and 1.24 for a total population of 1000, 1.31 for a population of 1500, 1.41 for a population of 2250, 1.51 for a population of 3000, 1.78 for a population of 5000, 2.95 for a population of 7000, and 6.96 for a population of 10000. Interestingly this did not necessarily vary monotonically as a



function of smart partial population, as was the case with latency.

% Smart	500	750	1000	1500	2250	3000	5000	7000	10000
0%-2.5%	null	null	null	null	null	null	null	null	null
2.5%-5%	0.3731	0.1128	0.3973	0.0302	0.2641	0.1992	0.6728	0.5995	1.2449
5%-7.5%	-0.594	0.3011	-0.0796	0.204	0.3692	0.3828	0.3722	1.0331	1.3481
7.5%-10%	0.5474	0.0772	0.0836	0.1319	-0.144	0.1149	0.4276	0.6898	5.251
10%-12.5%	0.0548	0.0742	-0.0258	-0.06	0.3468	0.2454	0.3595	1.1701	14.4703
12.5%-15%	-0.1514	0.1939	0.1033	0.2027	0.1692	0.1188	0.6015	0.4957	9.3593
15%-17.5%	-0.1222	-0.1011	0.2545	0.0996	0.0676	0.2683	0.2444	2.0921	11.8828
17.5%-20%	0.1335	0.1734	-0.0771	0.0914	0.2588	0.2086	0.392	5.9932	14.8143
20%-22.5%	0.0093	0.0115	0.1335	0.096	0.0762	0.1464	0.3459	5.4107	4.2493
22.5%-25%	0.0757	0.0923	0.006	0.1101	0.1341	0.152	1.6189	13.095	9.8487

Table 5.4: Percentage change in mean smart vehicle latency at 2.5% increments of smart partial population.

The distance ratio for smart drivers blows up at 17.5% in a population of 7000 and 12.5% in a population of 10000. At 25%, smart drivers suffer from a distance ratio greater than seven in a population of 7000, and greater than 13 for a population of 10000.

% Smart	500	750	1000	1500	2250	3000	5000	7000	10000
0%-2.5%	1.2061	1.2052	1.2391	1.3359	1.4144	1.5233	1.7582	2.0013	2.3645
2.5%-5%	1.1497	1.2302	1.2371	1.3118	1.4054	1.5149	1.737	2.002	2.3344
5%-7.5%	1.1739	1.2231	1.2495	1.3132	1.4081	1.5156	1.7602	1.9702	2.3645
7.5%-10%	1.1677	1.2117	1.2377	1.3116	1.4145	1.5052	1.7562	1.9893	2.953
10%-12.5%	1.1829	1.1976	1.2457	1.2991	1.4204	1.5009	1.743	1.9869	5.8039
12.5%-15%	1.1666	1.2189	1.2367	1.3113	1.4046	1.5018	1.7534	1.9848	7.3791
15%-17.5%	1.1758	1.222	1.251	1.3239	1.4101	1.4966	1.7552	2.2691	9.4225
17.5%-20%	1.1774	1.2146	1.241	1.3199	1.4012	1.499	1.7508	3.5007	11.6462
20%-22.5%	1.1767	1.2184	1.2481	1.3123	1.4061	1.5103	1.7497	4.6679	11.9725
22.5%-25%	1.1749	1.2096	1.2471	1.3098	1.4064	1.5003	2.0717	7.1645	13.3683

Table 5.5: Distance ratios for smart vehicles as a function of partial population.

Both speed and latency metrics indicate a performance roll-off and, ultimately, severe degradation beyond certain proportions of smart drivers in larger populations. This is likely attributable to constraints of the network itself. There exists a finite number of paths between any origin and destination. Any link that appears in that path may also appear in the paths of other OD-pairs. In cases of large commuter populations, actively rerouting

vehicles to avoid congestion is largely futile when most links are *already* highly congested. Congestion-conscious link selections are only marginally better than congestion-ignorant ones, and have the effect of keeping vehicles on the network far longer than they would be otherwise. As the number of smart drivers increases, one witnesses the emergence of competition between them; smart vehicles are aware of other smart vehicles' locations, but not their destinations or planned paths. This will cause them to recursively reroute to avoid one another, resulting in detrimental effects to their own driver class and stagnation of the performance of the system.

Practically speaking, the results indicate that in small and moderately-sized populations, rerouting can ease congestion and improve individual itineraries without incurring severe distance penalties. However, in large populations, the limiting factor simply becomes the amount of pavement on the network. Improvement resulting from rerouting will only occur when a small fraction of drivers practice such behavior in these populations. As that fraction increases, one observes the emergence of competitive behaviors between such drivers.

### 5.3 Model Calibration

The NetLogo model is calibrated according to metrics adapted from Homburger, Perkins, and Kell [50], adjusted to reflect real-world data documented and made available by the Virginia Department of Transportation (VDOT)<sup>1</sup>, the Texas Transportation Institute (TTI)<sup>2</sup>, and the Center for Advanced Transportation Technology Laboratory at the University of Maryland (CATT Lab)<sup>3</sup>. Establishing an upper performance bound of 10,000 vehicles/25% smart partial population, the model was calibrated so that the worst congestion conditions would occur at that total population (at an unknown smart partial population, hypothesized to be zero percent). This resulted in a minimum vehicle speed of 55 patches per tick (32.5 miles per hour, the approximate upper bound of level-of-service F). This was set as the

---

<sup>1</sup><http://www.virginiadot.org/info/ct-trafficcounts.asp>

<sup>2</sup><http://mobility.tamu.edu/ums/congestion-data/>

<sup>3</sup><http://www.cattlab.umd.edu/>

lower bound because modeling an unstable flow regime (i.e. LOS F) requires implementing stop-and-go movement mechanics, which was not practically feasible given the scope and scale of the model.

It is necessary to note that itineraries and road use are distributed isotropically; that is, any origin-destination pair is as likely as any other. This distribution is time-independent over the duration of the simulation. In reality, certain roads experience substantially more use than others (interstate highways experience higher traffic than VA state highways), and roads experience different levels of use at different times of the day (I-66E experiences more traffic in the morning as people commute into Washington, DC for work, and less in the evening as those commuters are returning home on I-66W).

The test populations represent levels of use assuming a given real-world sample scenario is characteristic of the entire test system at any given point in time. Again, while this is obviously not what is observed in the real-world, it is adequate and convenient as far as exploratory work is concerned. Real-world speed values were obtained from the Regional Integrated Transportation Information System (RITIS) at the University of Maryland's CATT Lab<sup>4</sup>. RITIS aggregates raw speed data from motorists using INRIX on enabled devices, and compiles usage statistics on intervals ranging from five minutes to one hour. Table 4.1 compares test populations are zero percent smart vehicles to weekday (Monday through Friday) conditions during the month of July 2013. This month was selected due to the absence of inclement weather; precipitation and snowfall can adversely affect driving conditions, introducing a complicating factor that would otherwise need to be controlled for. Values were based on fifteen minute averages over morning hours.

---

<sup>4</sup><https://www.ritis.org/>

Total Population	Mean Simulated Speed (ppt)	Mean Simulated Speed (mph)	Real-World World Example	at Time	Comparable Real-Observed Speed (mph)	Level of Service
500	94.82	56.03	I-66W at Exit 73	05:00 Weekday Average	56.01	A
750	92.13	54.44	I-66E at Exit 75	07:00 Weekday Average	54.71	B
1000	89.59	52.94	I-395S at Exit 10	08:45 Weekday Average	52.99	B
1500	84.89	50.16	I-66W at Exit 73	09:00 Weekday Average	50.61	C
2250	79.55	47.01	VA-110N at Memorial Drive	09:00 Weekday Average	47.39	C
3000	75.58	44.66	US-50E at Fort Myer Drive	07:45 Weekday Average	44.66	D
5000	68.8	40.65	US-50W at US-27	08:00 Weekday Average	40.05	D
7000	65.14	38.49	US-50E at N. Courthouse Road	08:00 Weekday Average	38.15	E
10000	61.49	36.34	US-29S at N Veitch St	08:00 Weekday Average	36.56	E/F

Table 5.6: Comparison between simulation speeds and observed speeds for weekday mornings, July 2013.

## 5.4 Verification and Validation

Verification and validation are inherently difficult to perform on agent-based models (when compared to other simulation methodologies), and there exists a substantial body of literature both describing this difficulty, as well as indicating the absence of consensus on how to do so [51–55]. However, it is possible to put forth an *ad hoc* approach for this specific case. In order to employ this methodology in simulating ATIS use and road network congestion, it is necessary to understand that results are bounded by the assumptions and mechanics of the simulation itself. First, driver behaviors and preferences are, in reality, often opaque, meaning that decision-making criteria is not necessarily obvious to a third-party and is often not consistently followed by the drivers themselves. This means classifying drivers on a set of very straightforward routing criteria (as done here) that are assumed to be consistently followed will produce constrained results. Second, the largest population tested was 10000 vehicles, whereas the average annual daily traffic of the highest-volume sample segments numbers in the tens of thousands<sup>5</sup>. Modeling populations of this size is computationally impractical, especially using a platform like NetLogo. Third, the origins and destinations of vehicles was generated randomly and transit corridors were treated as isotropic; this does not appropriately reflect a workday commuter scenario, where certain origins, destinations and thoroughfares are more likely to be visited than others at some point in the day. Fourth, trading granularity for scope

<sup>5</sup>From [virginiadot.org/info/2013\\_traffic\\_data.asp](http://virginiadot.org/info/2013_traffic_data.asp). A section of I-66 in Fairfax County had an AADT of 96,000 in 2013.

results in the omission of certain factors that do impact congestion. Acceleration/deceleration, car-following, and lane-changing behaviors all critically impact traffic flow and the likelihood of a vehicle accident. Infrastructure characteristics, such as lane width, lane merging/splitting, and turn curvature, are also important in flow modeling and accident prediction. However, as indicated in the literature review, these are appropriately modeled on a much smaller scope by means of microsimulation, and would be impractical to model in a simulation of this scale and scope.

Validation is complicated by the fact that data pertaining to ATIS provided by commercial entities is not publicly available for analytics purposes. Organizations that offer free ATIS services, notably Google Maps, Waze, Apple Maps, Nokia HERE, and INRIX do not make their navigation data or user statistics openly available. Google does collect location data by means of cellular and aGPS tracking on users who have agreed to this service, even for users who are not currently using Maps or another guidance platform [56]. However, the number of Northern Virginia commuters who use ATIS-enabled phones or devices, have location services enabled, and at any given point are using Google Maps for route guidance purposes, is not public information. It should be noted, however, that the cited article suggests a hypothetical scenario in which 5% of Google Maps users could be actively rerouted, indicating that, at the time of its authorship, the actual amount is far lower.

The results here are more powerful to the effect of informing high-level policy and implementation targets for said services. That is, the documented methodology and outcomes are more valuable as a methodological proof of concept than they are a highly-accurate predictive tool. However, it is possible to present a validation procedure that could be implemented in future work. Validation would consist of long-term observation of a real-world traffic network (in this case, the used portion of Northern Virginia), with a high degree of information on who and how many are using ATIS guidance, and when and under what circumstances drivers are using them. Such an effort would require cooperation, coordination, and some amount of participatory modeling between consumer and commercial

ATIS providers, public safety and law enforcement officers and officials, local and state government, and regular motorists (who would need to consent to being surveyed). This would allow for better constraint revisions and calibration; moreover, it would inform whether emergent outcomes are practically and politically actionable and/or desirable. This is further explored in sections 6.3 and 6.4.

# Conclusions and Considerations, Future Work, and Applications

## 6.1 Conclusions and Contributions

Not surprisingly, rerouting a fraction of agents did result in, in most cases, a modest improvement of mean vehicle speed and in latency relative to the theoretical maximum across the entirety of the network and all agents on it. This improvement was most substantial in medium and medium-large vehicle populations. However, very large populations experienced the emergence of competitive pressure between constantly-rerouting vehicles. This suggests that further development of realtime rerouting by ATIS platforms and continued adoption by motorists should have beneficial effects on easing systemic congestion in high-traffic metro areas. However, it is also apparent that the network itself did have a bounding effect on the potential for improvement by way of this approach. While agent-based modeling has been successfully used in the past to evaluate traffic augmentation by way of ATIS, this research demonstrates the efficacy of applying it to a larger area, more complex network, and different methodological approach (mesoscopic instead of microscopic) than has appeared in the traffic engineering literature.

## 6.2 Considerations

While the improvements may seem marginal, the cumulative effect when applied to tens or hundreds of thousands commuters is nontrivial. Recall that the *Urban Mobility Report* lists the economic cost of congestion at \$121bn. An adoption of 2.5% ATIS use and compliance with a total improvement of latency by 0.4% would therefore reduce this loss by \$500M. It is

therefore necessary to keep in mind the aggregate potential of ATIS implementation, even if the individual effects seem trivial.

## 6.3 Future Work

Future work should include supplementing the simulation with both accurate driver survey and demographic data to better inform driver models, as well as road sensor data to better represent typical commutes as a function of time of day. Most importantly, it should include data, wherever available, from consumer ATIS providers that describes how drivers respond to presented data. This would allow for the establishment of a better experimental baseline, preallocating agents to specific roadways and OD-pairs that are most likely to experience traffic. This would also enable more sophisticated preference coding for the agents themselves.

Additionally, it would be advisable to test on another geospatial dataset; while it is intuitively reasonable to assume that a traffic easing methodology that does not directly rely on local network topology will work on any nontrivial network, it is likely that results will vary between locations because of network connectivity and infrastructure assets, and certain metro areas will enjoy greater benefit than others using this approach. Given appropriate geospatial data, the existing simulation can easily be used to test somewhere other than the Northern Virginia area investigated here.

Finally, it is suggested that structural components be augmented and competitive pressures between them implemented. This simulation assumes that there exists a single entity that collects, stores, and disseminates the state of the network, and that vehicles are routed and rerouted accordingly. It may be beneficial to test how competing ATIS providers affect the network (for example, some users use Google Maps, others use Apple Maps, others use Nokia HERE). Granting privileges to a local “Department of Transportation” agent that can affect the accessibility of the network would also be important to investigate. For example, how does the presence of HOV lanes, traffic rerouting by means of policing, and congestion value



pricing affect commuter travel and possibly conflict with stated ATIS routes? These are all features to consider when extending the research presented here for more comprehensive system modeling.

## 6.4 Applications

The most obvious applications of the described research pertain to planning, testing, and generalized prediction of traffic and congestion as a function of ATIS implementation. First, this approach could be used to inform ATIS adoption targets by drivers for mandated levels of congestion improvement. Conversely, it can be used to predict how much congestion will improve given an observed or hypothetical level of ATIS use. This is potentially useful for consumer navigation services, such as Google Maps, Waze, Nokia HERE, and INRIX, as they continue to roll out rerouting features on their respective mobile phone applications.

Applications are also promising for municipalities and transportation agencies that, as stated in the introduction and problem overview, currently suffer from (and further face) extensive overuse of roadway infrastructure. This methodology would supplement other strategies in a multimodal approach. Combining congestion easing by means of ATIS with incentives, such as HOV lanes and congestion value pricing, has the potential to reduce systemic congestion on high-traffic corridors.

This simulation could be retooled for industry-specific applications, as well. ABM's primary appeal in the modeling field is in its ability to observe how individual entities, with some degree of intelligence, interact with one another and their environment. Employing this simulation with modified agent types or the addition of new ones would be promising for industries such as commercial shipping and logistics, whose business can be impacted by congestion-related delays.

# Bibliography

- [1] David Schrank, Tim Lomax, and Bill Eisele. 2011 urban mobility report. 2011.
- [2] Erik Hansson, Kristoffer Mattisson, Jonas Bjork, Per-Olof Ostergren, and Kristina Jakobsson. Relationship between commuting and health outcomes in a cross-sectional population survey in southern sweden. *BMC Public Health*, 11(1):834, 2011.
- [3] Tim Lomax, David Schrank, Shawn Turner, and Richard Margiotta. Selecting travel reliability measures. *Texas Transportation Institute monograph (May 2003)*, 2003.
- [4] Paul Davidsson, Lawrence Henesey, Linda Ramstedt, Johanna Tornquist, and Fredrik Wernstedt. An analysis of agent-based approaches to transport logistics. *Transportation Research part C: emerging technologies*, 13(4):255–271, 2005.
- [5] Dirk Helbing. Social self-organization. *Social Self-Organization: Agent-Based Simulations and Experiments to Study Emergent Social Behavior, Understanding Complex Systems*. ISBN 978-3-642-24003-4. Springer-Verlag Berlin Heidelberg, 2012, 1, 2012.
- [6] Thomas Schulze and Thomas Fliess. Urban traffic simulation with psycho-physical vehicle-following models. In *Proceedings of the 29th conference on Winter simulation*, pages 1222–1229. IEEE Computer Society, 1997.
- [7] Feirouz Ksontini, Rene Mandiau, Zahia Guessoum, and Stephane Espie. Affordance-based agent model for road traffic simulation. *Autonomous Agents and Multi-Agent Systems*, pages 1–29, 2014.
- [8] Laron Smith, Richard Beckman, Doug Anson, Kai Nagel, and Michael E Williams. Transims: Transportation analysis and simulation system. In *Fifth National Conference on Transportation Planning Methods Applications-Volume II: A Compendium of Papers Based on a Conference Held in Seattle, Washington in April 1995*, 1995.
- [9] Massimiliano Daniele Rosini. Macroscopic models for vehicular flows and crowd dynamics: Theory and applications. *Macroscopic Models for Vehicular Flows and Crowd Dynamics: Theory and Applications: Classical and Non-Classical Advanced Mathematics for Real Life Applications, Understanding Complex Systems*. ISBN 978-3-319-00154-8. Springer International Publishing, 1, 2013.
- [10] Nicola Bellomo, M. Delitala, and V. Coscia. On the mathematical theory of vehicular traffic flow i: Fluid dynamic and kinetic modelling. *Mathematical Models and Methods in Applied Sciences*, 12(12):1801–1843, 2002.

- [11] Wilco Burghout, Haris N. Koutsopoulos, and Ingmar Andreasson. A discrete-event mesoscopic traffic simulation model for hybrid traffic simulation. In *Intelligent Transportation Systems Conference, 2006. ITSC'06. IEEE*, pages 1102–1107. IEEE, 2006.
- [12] Ed Manley, Tao Cheng, Alan Penn, and Andy Emmonds. A framework for simulating large-scale complex urban traffic dynamics through hybrid agent-based modelling. *Computers, Environment and Urban Systems*, 44:27–36, 2014.
- [13] Wilco Burghout, Haris N. Koutsopoulos, and Ingmar Andreasson. Hybrid mesoscopic-microscopic traffic simulation. *Transportation Research Record: Journal of the Transportation Research Board*, 1934(1):218–255, 2005.
- [14] Ye Tian and Yi-Chang Chiu. Anisotropic mesoscopic traffic simulation approach to support large-scale traffic and logistic modeling and analysis. In *Proceedings of the Winter Simulation Conference*, pages 1500–1512. Winter Simulation Conference, 2011.
- [15] Michael J Lighthill and Gerald Beresford Whitham. On kinematic waves. ii. a theory of traffic flow on long crowded roads. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 229(1178):317–345, 1955.
- [16] Paul I Richards. Shock waves on the highway. *Operations research*, 4(1):42–51, 1956.
- [17] Harold J Payne. Models of freeway traffic and control. *Mathematical models of public systems*, 1971.
- [18] GB Whitham. Linear and nonlinear waves. 1974.
- [19] George Rowlands, E Infeld, and AA Skorupski. Some exact solutions to the lighthill–whitham–richards–payne traffic flow equations. *Journal of Physics A: Mathematical and Theoretical*, 46(36):365202, 2013.
- [20] Kai Nagel and Michael Schreckenberg. A cellular automaton model for freeway traffic. *Journal de physique I*, 2(12):2221–2229, 1992.
- [21] Marcus Rickert, Kai Nagel, Michael Schreckenberg, and Andreas Latour. Two lane traffic simulations using cellular automata. *Physica A: Statistical Mechanics and its Applications*, 231(4):534–550, 1996.
- [22] Peter Wagner, Kai Nagel, and Dietrich E. Wolf. Realistic multi-lane traffic rules for cellular automata. *Physica A: Statistical Mechanics and its Applications*, 234(3):687–698, 1997.
- [23] Partich AM Ehlert and Leon JM Rothkrantz. Microscopic traffic simulation with reactive driving agents. In *Intelligent Transportation Systems, 2001. Proceedings. 2001 IEEE*, pages 860–865. IEEE, 2001.
- [24] Hong Zheng, Young-Jun Son, Yi-Chang Chiu, Larry Head, Yiheng Feng, Hui Xi, So-jung Kim, and Mark Hickman. *A Primer for Agent-Based Simulation and Modeling in Transportation Applications*, 2013.

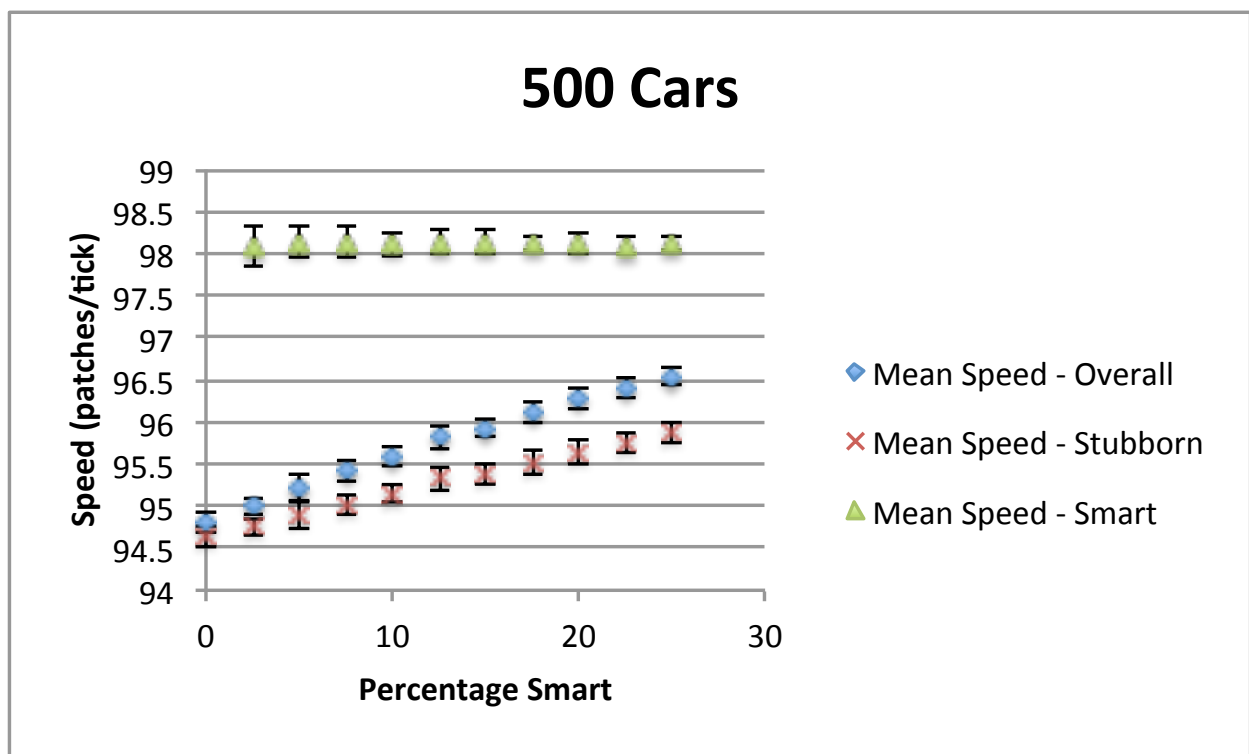
- [25] K. Bernhardt. Agent-based modeling in transportation. *Artificial Intelligence in Transportation*, 72, 2007.
- [26] Hussein Dia. An agent-based approach to modelling driver route choice behaviour under the influence of real-time information. *Transportation Research Part C: Emerging Technologies*, 10(5):331–349, 2002.
- [27] Joachim Wahle, Ana Lucia C. Bazzan, Franziska Klugl, and Michael Schreckenberg. The impact of real-time information in a two-route scenario using agent-based simulation. *Transportation Research Part C: Emerging Technologies*, 10(5):399–417, 2002.
- [28] Joachim Wahle, Ana Lucia C. Bazzan, Franziska Klugl, and Michael Schreckenberg. Decision dynamics in a traffic scenario. *Physica A: Statistical Mechanics and its Applications*, 287(3):669–681, 2000.
- [29] Daniele Buscema, Matteo Ignaccolo, Giuseppe Inturri, Alessandro Pluchino, Andrea Rapisarda, Corrado Santoro, and Salvatore Tudisco. The impact of real time information on transport network routing through intelligent agent-based simulation. In *Science and Technology for Humanity (TIC-STH), 2009 IEEE Toronto International Conference*, pages 72–77. IEEE, 2009.
- [30] Michael Balmer, Nurhan Cetin, Kai Nagel, and Bryan Raney. Towards truly agent-based traffic and mobility simulations. In *Autonomous Agents and Multiagent Systems, 2004. AAMAS 2004. Proceedings of the Third International Joint Conference on*, pages 60–67. IEEE, 2004.
- [31] Reinhart D. Kuhne. Foundations of traffic flow theory i: Greenshields’ legacy: Highway traffic. In *Symposium on the Fundamental Diagram: 75 Years (Greenshields 75 Symposium)*, 2008.
- [32] Frank A. Haight. *Mathematical theories of traffic flow*. 1965.
- [33] John Glen Wardrop. Road paper. some theoretical aspects of road traffic research. In *ICE Proceedings: Engineering Divisions*, volume 1, pages 325–362. Thomas Telford, 1952.
- [34] Martin Beckmann, CB McGuire, and Christopher B. Winsten. *Studies in the Economics of Transportation*, 1956.
- [35] John N. Tsitsiklis and Dimitri P. Bertsekas. Distributed asynchronous optimal routing in data networks. *Automatic Control, IEEE Transactions on*, 31(4):325–332, 1986.
- [36] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *STACS 99*, pages 404–413. Springer, 1999.
- [37] Tim Roughgarden and Ava Tardos. How bad is selfish routing? *Journal of the ACM (JACM)*, 49(2):236–259, 2002.

- [38] Yannis A. Korilis, Aurel A. Lazar, and Ariel Orda. Avoiding the braess paradox in non-cooperative networks. *Journal of Applied Probability*, 36(1):211–222, 1999.
- [39] Richard Steinberg and Willard I. Zangwill. The prevalence of braess’ paradox. *Transportation Science*, 17(3):301–318, 1983.
- [40] Christos Papadimitriou and Gregory Valiant. Selfish routers and the price of anarchy.
- [41] Tim Roughgarden. Selfish routing and the price of anarchy. 2005.
- [42] Baruch Awerbuch, Yossi Azar, and Amir Epstein. The price of routing unsplittable flow. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 57–66. ACM, 2005.
- [43] George Christodoulou and Elias Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 67–73. ACM, 2005.
- [44] Renato Paes Leme and Eva Tardos. Pure and bayes-nash price of anarchy for generalized second price auction. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 735–744. IEEE, 2010.
- [45] Tim Roughgarden. The price of anarchy in games of incomplete information. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 862–879. ACM, 2012.
- [46] Seraphim Alvanides and Godwin Yeboah. Cycling in the city: Understanding urban cycling behaviours and constraints in space-time. 2012.
- [47] Bin Jaing. Axwoman 6.0: Movingbehavior1 and movingbehavior2.
- [48] Uri Wilensky. Link-walking turtles, 2013.
- [49] Uri Wilensky. Move towards target, 2013.
- [50] WS Homburger, JH Kell, and D Perkins. Highway traffic noise: fundamentals of traffic engineering, 1992.
- [51] Paul Ormerod and Bridget Rosewell. Validation and verification of agent-based models in the social sciences. In *Epistemological Aspects of Computer Simulation in the Social Sciences*, pages 130–140. Springer, 2009.
- [52] David Midgley, Robert Marks, and Dinesh Kunchamwar. Building and assurance of agent-based models: An example and challenge to the field. *Journal of Business Research*, 60(8):884–893, 2007.
- [53] Xiaorong Xiang, Ryan Kennedy, Gregory Madey, and Steve Cabaniss. Verification and validation of agent-based scientific simulation models. In *Agent-Directed Simulation Conference*, pages 47–55, 2005.

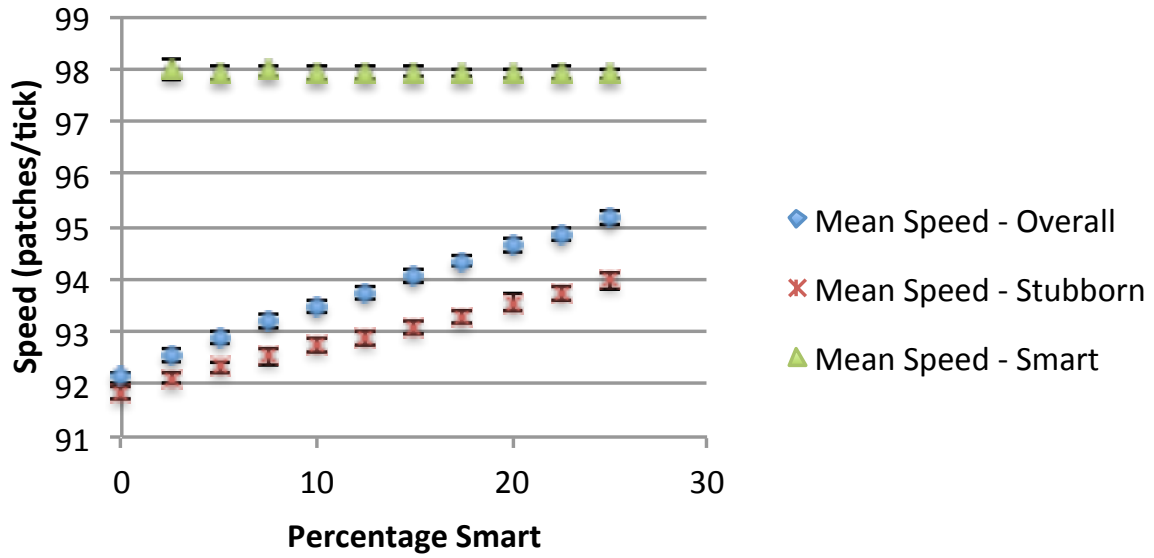
- [54] Uri Wilensky and William Rand. Making models match: Replicating an agent-based model. *Journal of Artificial Societies and Social Simulation*, 10(4):2, 2007.
- [55] Paul Windrum, Giorgio Fagiolo, and Alessio Moneta. Empirical validation of agent-based models: Alternatives and prospects. *Journal of Artificial Societies and Social Simulation*, 10(2):8, 2007.
- [56] Susan E. Matthews. How google tracks traffic, 03 July 2013.
- [57] V. Ljubovic. Traffic simulation using agent-based models. In *Information, Communication and Automation Technologies, 2009. ICAT 2009. XXII International Symposium on*, pages 1–6. IEEE, 2009.
- [58] LH Immers and S Logghe. Traffic flow theory. *Faculty of Engineering, Department of Civil Engineering, Section Traffic and Infrastructure, Kasteelpark Arenberg, 40*, 2002.

# Appendix A


Below are the individual speed plots indicated in Figure 5.1. Units are scaled such that a speed of 110 corresponds to 65 miles per hour. Bars indicate one standard deviation of the mean.



## 750 Cars




## 1000 Cars

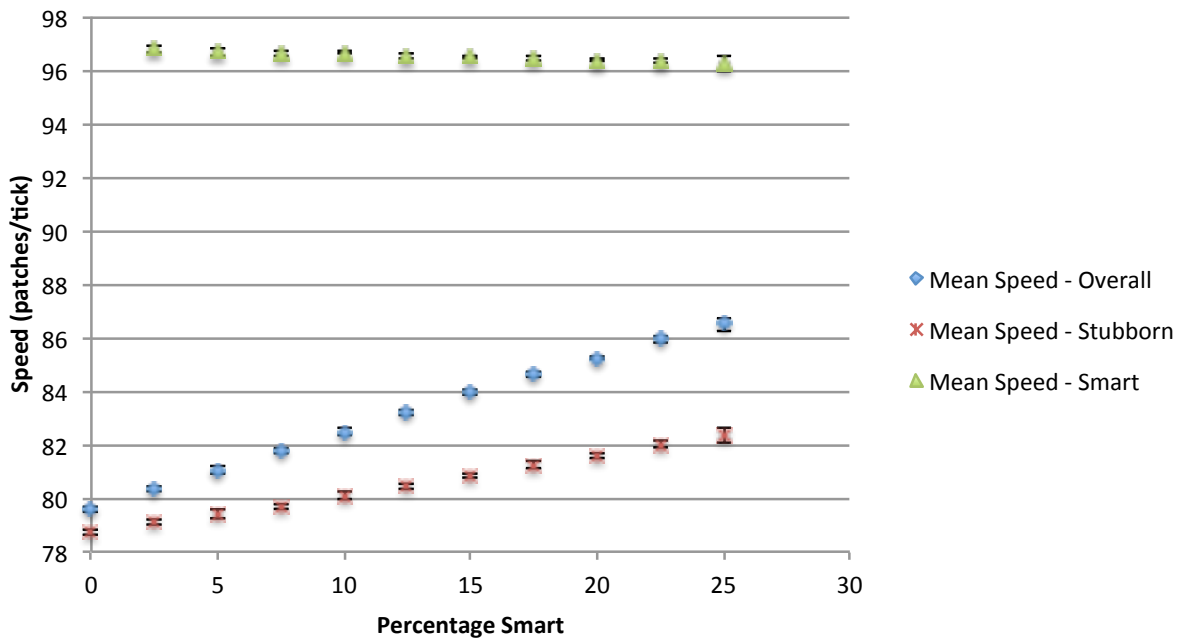


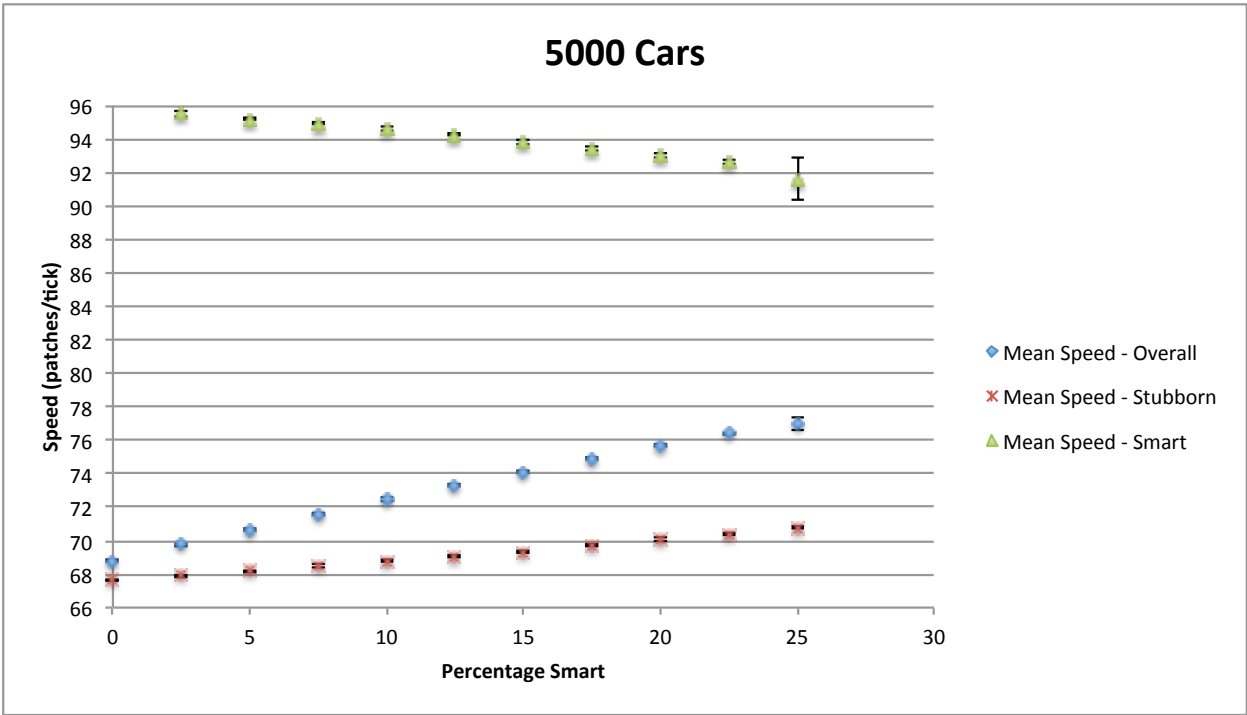
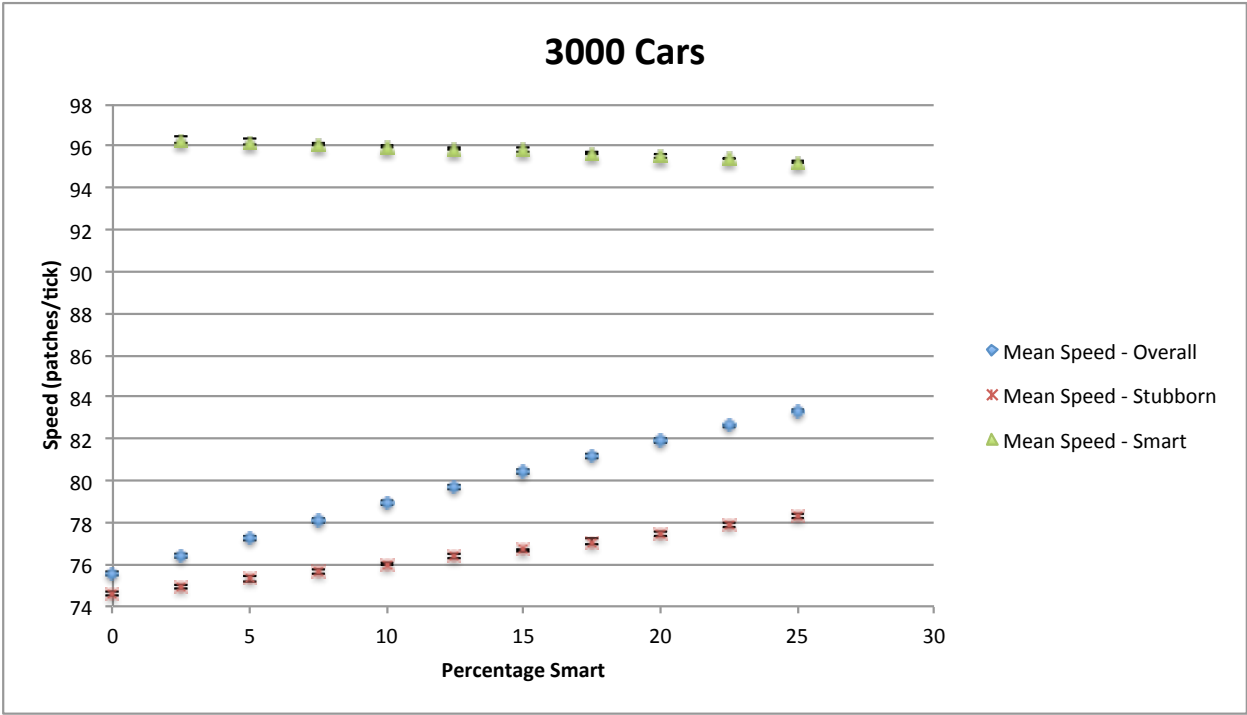



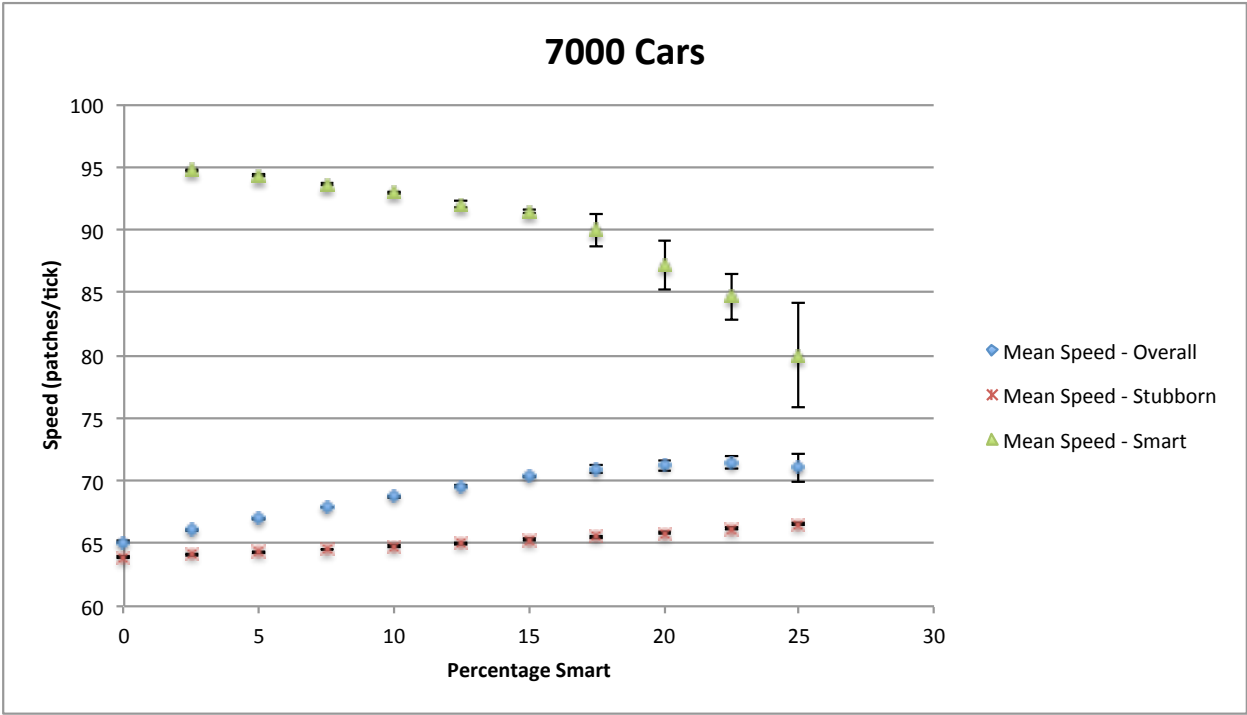
# 1500 Cars



# 2250 Cars







# Appendix B

Below are tables indicated the percentage change in speed for stubborn and smart vehicles.

% Smart	500	750	1000	1500	2250	3000	5000	7000	10000
0%-2.5%	0.1203	0.2985	0.2259	0.3763	0.4669	0.4522	0.3815	0.3264	0.2765
2.5%-5%	0.1442	0.2268	0.3673	0.4084	0.3758	0.4529	0.4043	0.2884	0.3223
5%-7.5%	0.1401	0.2194	0.2259	0.3485	0.3827	0.4765	0.4616	0.3636	0.3448
7.5%-10%	0.133	0.2102	0.2602	0.3055	0.4603	0.4853	0.3855	0.3388	0.3946
10%-12.5%	0.1765	0.1648	0.3262	0.42	0.4545	0.4725	0.4459	0.3751	0.4256
12.5%-15%	0.0623	0.229	0.2614	0.4306	0.541	0.397	0.3738	0.4361	0.494
15%-17.5%	0.1494	0.2294	0.3441	0.4029	0.4693	0.5217	0.5382	0.3664	0.5272
17.5%-20%	0.1276	0.2656	0.3256	0.4109	0.4334	0.5064	0.5191	0.4908	0.5341
20%-22.5%	0.1223	0.1902	0.3083	0.4029	0.509	0.4978	0.4691	0.4769	0.4564
22.5%-25%	0.117	0.2569	0.2277	0.4461	0.4293	0.5926	0.5827	0.6165	0.6192

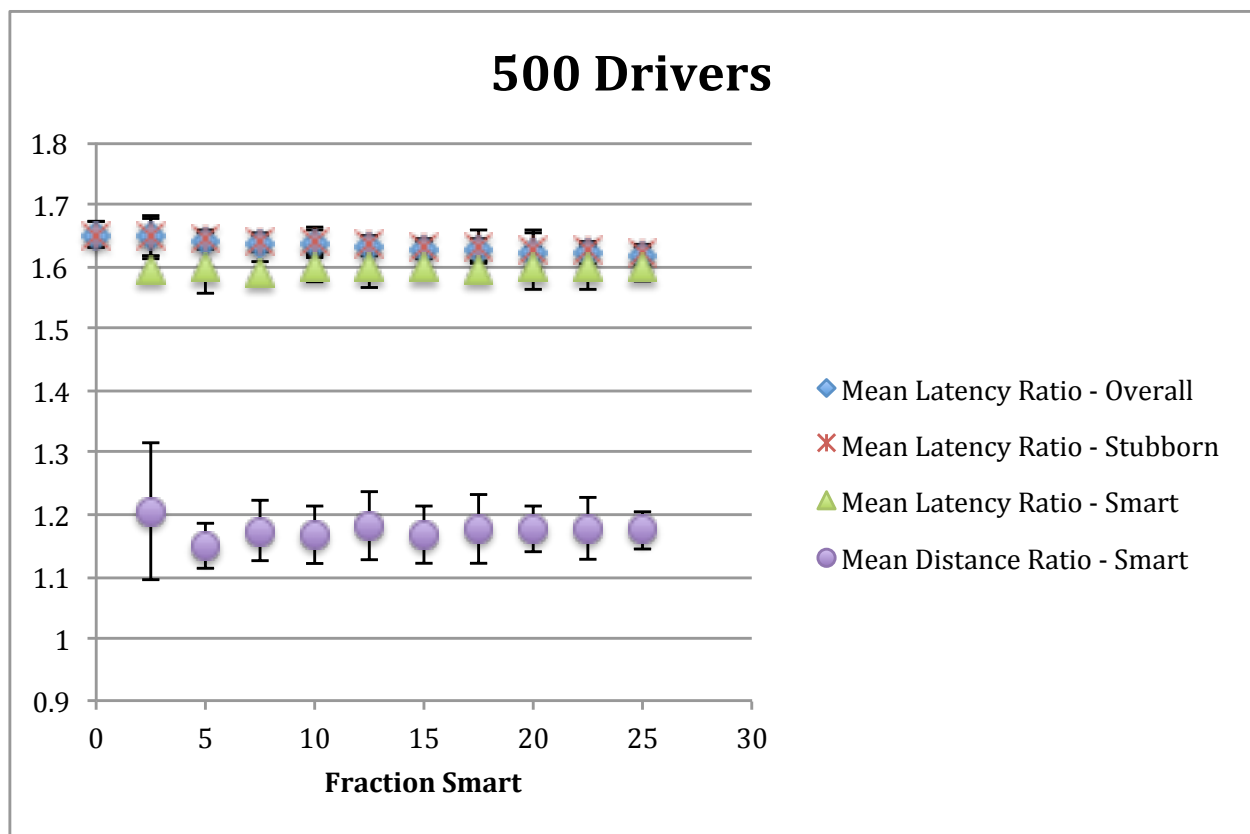
Table 1: Relative percentage change in speed for stubborn vehicles as a function of the fraction of smart vehicles.

% Smart	500	750	1000	1500	2250	3000	5000	7000	10000
0%-2.5%	null	null	null	null	null	null	null	null	null
2.5%-5%	0.0347	-0.0467	-0.0699	0.0524	-0.0825	-0.0744	-0.3055	-0.498	-1.0126
5%-7.5%	0.0046	0.0346	-0.0388	-0.0227	-0.072	-0.1227	-0.2777	-0.7156	-1.3406
7.5%-10%	-0.0242	-0.0319	-0.0206	-0.0039	0.0402	-0.0944	-0.3662	-0.6925	-3.2052
10%-12.5%	0.0237	-0.0014	0.0445	0.014	-0.1598	-0.1107	-0.3486	-1.0182	-7.0823
12.5%-15%	-0.0073	0.0137	0.0092	-0.0471	-0.0129	-0.091	-0.4888	-0.6319	-4.1397
15%-17.5%	-0.0029	-0.0119	-0.0412	0.0064	-0.045	-0.1382	-0.412	-1.6276	-4.7145
17.5%-20%	-0.0089	-0.0083	0.0254	-0.0247	-0.0911	-0.1288	-0.4436	-3.1149	-5.4319
20%-22.5%	-0.017	0.0001	-0.0357	-0.0164	-0.0112	-0.1373	-0.3893	-2.9002	-1.1018
22.5%-25%	0.009	-0.0118	-0.0023	-0.0288	-0.1307	-0.1696	-1.102	-5.4641	-3.4391

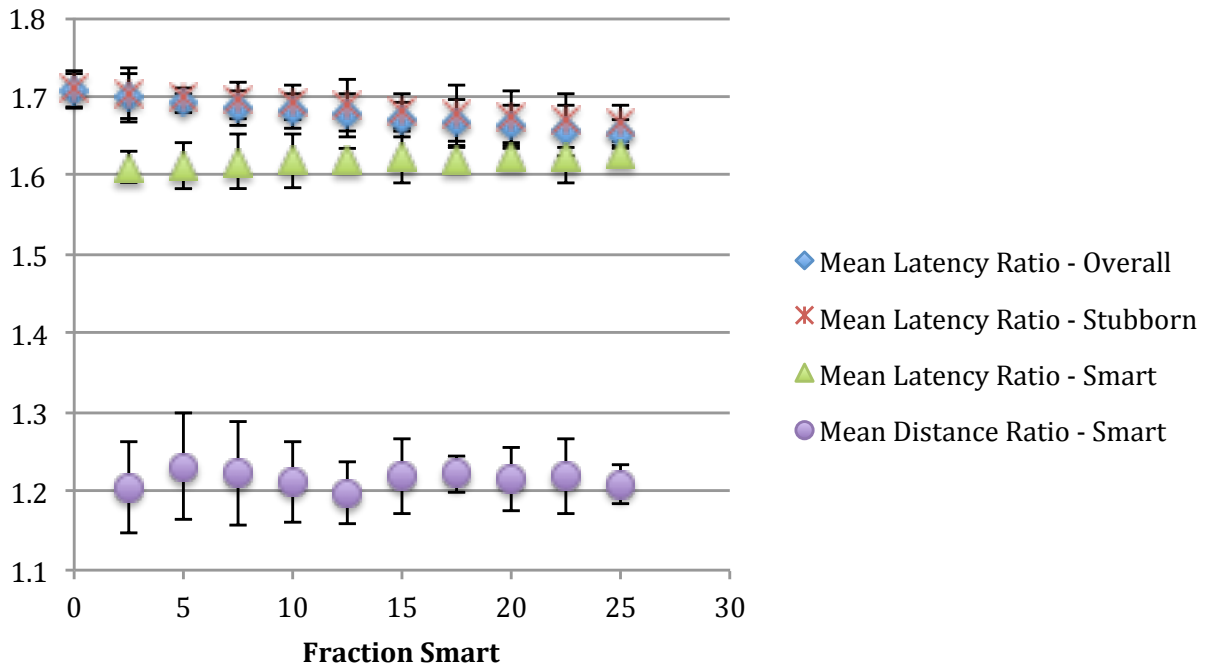
Table 2: Relative percentage change in speed for smart vehicles as a function of the fraction of smart vehicles.

# Appendix C


Below are the individual latency plots indicated in Figure 5.3. Units are scaled such that a speed of 110 corresponds to 65 miles per hour. Bars indicate one standard deviation of the mean.

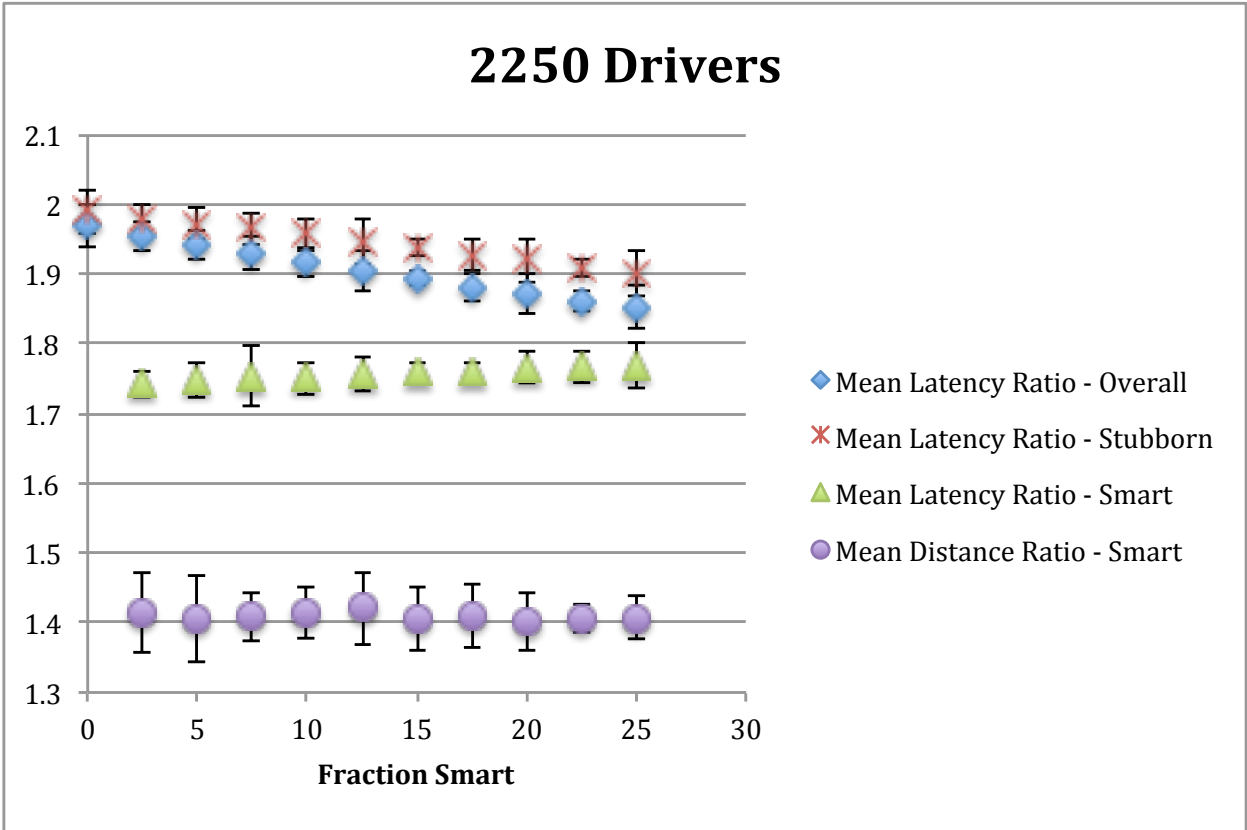



## 750 Drivers

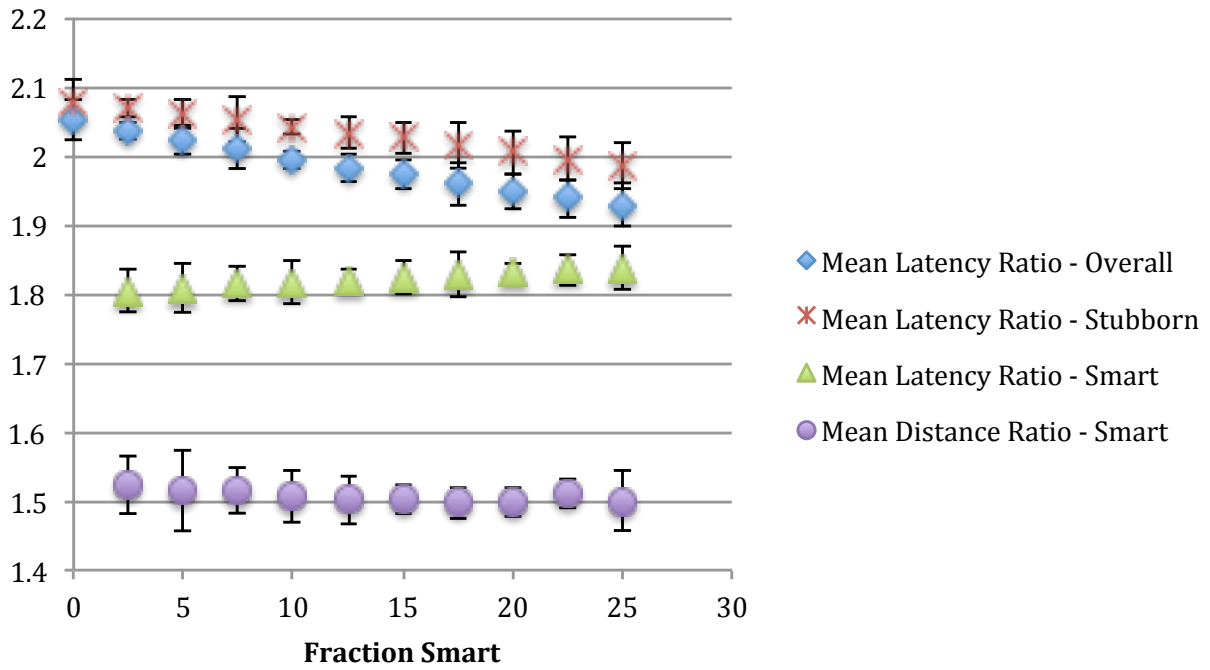


## 1000 Drivers





### 3000 Drivers



### 5000 Drivers

