

# Altruistic Autonomy: Beating Congestion on Shared Roads

Erdem Bıyık<sup>\*1</sup>, Daniel A. Lazar<sup>\*2</sup>, Ramtin Pedarsani<sup>2</sup>, Dorsa Sadigh<sup>3</sup>

<sup>1</sup> Electrical Engineering, Stanford University,

<sup>2</sup> Electrical and Computer Engineering, University of California, Santa Barbara,

<sup>3</sup> Computer Science & Electrical Engineering, Stanford University.

ebyik@stanford.edu, {dlazar, ramtin}@ece.ucsb.edu, dorsa@cs.stanford.edu

**Abstract.** Traffic congestion has large economic and social costs. The introduction of autonomous vehicles can potentially reduce this congestion, both by increasing network throughput and by enabling a social planner to incentivize users of autonomous vehicles to take longer routes that can alleviate congestion on more direct roads. We formalize these effects of *altruistic autonomy* on roads shared between human drivers and autonomous vehicles. In this work, we develop a formal model of road congestion on shared roads based on the fundamental diagram of traffic. We consider a network of parallel roads and provide algorithms that compute optimal equilibria that are robust to additional unforeseen demand. We further plan for optimal routings when users have varying degrees of altruism. We find that even with arbitrarily small altruism, total latency can be unboundedly better than without altruism, and that the best selfish equilibrium can be unboundedly better than the worst selfish equilibrium. We validate our theoretical results through microscopic traffic simulations and show average latency decrease of a factor of 4 from worst-case selfish equilibrium to the optimal equilibrium when autonomous vehicles are altruistic.

**Keywords:** Autonomy-Enabled Transportation Systems, Multi-robot systems, Stackelberg Routing Game, Optimization

## 1 Introduction

Autonomous and connected vehicles are soon becoming a significant part of roads normally used by human drivers. Such vehicles hold the promise of safer streets, better fuel efficiency, more flexibility in tailoring to specific drivers' needs, and time savings. In addition to benefits impacting individual users, autonomous cars can significantly influence the flow of traffic networks. A coalition of autonomous cars can potentially decrease road congestion through a number of creative techniques. Autonomous vehicles can increase road capacity and throughput, as well

---

\* First two authors contributed equally and are listed in alphabetical order.

as decrease fuel consumption, by forming platoons and maintaining shorter headways to the vehicles they follow [1–5]. They can smooth shockwaves and instabilities that form in congested flows of human-driven vehicles [6–9]. Finally, autonomous vehicles can incentivize people to take *unselfish* routes, i.e., routes leading to a lower overall traffic delay for all drivers on the network. Previous work has studied methods such as tolling and incentives to influence drivers’ routing choices by, e.g., shifting their commute times or incentivizing them to use less congested roads or less congesting modes of transportation [10–12]. Our key insight is how people, who are incentivized to behave altruistically, along with vehicles with autonomous capabilities, can be leveraged to positively influence traffic networks. Our key contribution is to formally study and plan for *altruistic* autonomous cars that can influence traffic flows in order to achieve lower latencies. The setting we are interested in is as follows: assume there is a population of selfish human drivers and a population of autonomous car users, some of whom are willing to take routes with varying degrees of inefficiency. *How can we best utilize this altruistic autonomy to minimize average user delay?*

To answer this, we first develop a model of vehicle flow that reflects changes in road capacity due to efficient platooning of autonomous cars on shared roads. The latency experienced by users of a road is a function of the total vehicle flow as well as the *autonomy level*, the fraction of vehicles that are autonomous. We use the terminology *latency* to refer to the delay experienced by a driver from taking a road, *flow* to refer to the vehicle flux on a road (vehicles per second), and *density* to describe the number of cars per meter on a road. Our mixed-autonomy model is developed based on the commonly used fundamental diagram of traffic [13], which relates vehicle density to vehicle flow on roads with only human-driven cars.

Using our mixed-autonomy traffic flow model for a network of parallel roads, we study selfish equilibria in which human-drivers and users of autonomous vehicles all pick their routes to minimize their individual delays. As it may be difficult to exactly gauge flow demands, we define a notion of *robustness* that quantifies how resilient an equilibrium is to additional unforeseen demand. We establish properties of these equilibria and provide a polynomial-time algorithm for finding a robust equilibrium that minimizes overall delay.

We then address how to best use autonomous vehicles that are altruistic. We develop the notion of an *altruism profile*, which reflects the varying degrees to which users of autonomous vehicles are willing to take routes longer than the quickest route. We establish properties of optimal altruistic equilibria and provide a polynomial-time algorithm for finding such routings.

We validate our theoretical results through experiments in a microscopic traffic simulator. We show that an optimal selfish equilibrium can improve latency by a factor of two, and utilizing altruism can improve latency by another factor of two in a realistic driving scenario. Our contributions in this work are as follows:

- We develop a formal model of traffic flow on mixed-autonomy roads.

- We design an optimization-based algorithm to find a best-case Nash equilibrium.
- We define a robustness measure for the ability to allocate additional flow demand and develop an algorithm to optimize for this measure.
- We develop an algorithm to minimize total experienced latency in the presence of altruistic autonomous vehicles.<sup>4</sup>

Our development considers how to find optimal equilibria. This does not directly inform a planner on how to achieve such equilibria, but is the first step towards doing so.

**Related work.** Our work is closely related to work in optimal routing, traffic equilibria, and Stackelberg strategies in routing games. In routing games, populations of users travel from source graph nodes to destination nodes via edges, and the cost incurred due to traveling on an edge is a function of the volume of traffic flow on that edge. Some works have found optimal routing strategies for these games [15, 16]. If all users choose their routes selfishly and minimize their own travel time this results in a *Nash* or *Wardrop Equilibrium* [17, 18]. Stackelberg routing games are games in which a leader controls some fraction of traffic, and the remainder of the traffic reaches a selfish equilibrium [19–21]. Many works bound the inefficiency due to selfish routing in these games including in the presence of traffic composed of multiple vehicle types [22–25]. However, these works on routing games consider a model of road latency in which latency increases as vehicle flow increases. This latency model does not capture congestion, in which jammed traffic has low flow *and* high latency.

While not in the routing game setting, other works focus specifically on characterizing these congestion effects on traffic flow with multiple classes of users [26–29]. Krichene et al. combine these two fields by considering optimal Stackelberg strategies on networks of parallel roads with latency functions that model this notion of congestion [30]. Their work considers a single traffic type (no autonomous vehicles) and the controlled traffic can be made to take routes with delay arbitrarily worse than the best available route. In contrast, we consider a routing game with two types of traffic and there is a distribution of the altruism level of the controlled autonomous traffic.

## 2 Model and Problem Statement

**Traffic Flow on Pure Roads.** We assume every road has a maximum flow  $z_i^{\max}$ . This occurs when traffic is in *free-flow* – when all vehicles travel at the nominal road speed.

*Traffic density vs traffic flow:* Adding more cars to a road that is already at maximum flow makes the traffic switch from free-flow to a congested regime, which decreases the vehicle flow. In the extreme case, at a certain density, cars are bumper-to-bumper and vehicle flow stops. Removing vehicles also decreases the flow since it is not going to speed up cars that are already traveling in free-flow at their nominal speed. Fig. 1(a) – Fundamental Diagram of Traffic [13] – illustrates this phenomenon, where flow increases linearly with respect to density

---

<sup>4</sup> Due to space constraints, we defer all proofs to this paper’s extended version [14].

with slope  $v_i^f$  until it hits the critical density. The slope corresponds to the free-flow velocity on road  $i$ . At the critical point, flow decreases linearly until it is zero at the maximum density. This linear model matches our simulations (Section 4).




Fig. 1: (a) Fundamental diagram of traffic for traffic of a pure road, relating vehicle density and flow. (b) Relationship between vehicle flow and latency. Blue and brown lines correspond to free-flow and congested regimes; arrowheads show increasing density.

*Traffic flow vs road latency:* The relationship between vehicle flow and road latency reflects the same free-flow/congested divide. In free-flow the road has constant latency, as all vehicles are traveling at the nominal velocity. In the congested regime, however, vehicle flow decreases as latency increases – in congestion, a high density of vehicles is required to achieve a low traffic flow, resulting in high latency. This is represented in Fig. 1 (b), where in free-flow, latency is constant at free-flow latency  $a_i$  for any amount of flow up to the maximum flow on a road. The brown curve above corresponds to the congested regime, in which latency *increases* as vehicle flow decreases.<sup>5</sup>

**Traffic Flow on Mixed-Autonomy Roads.** We assume that on mixed-autonomy roads, the autonomous vehicles can coordinate with one another and potentially form platoons to help with the efficiency of the road network. We now extend the traffic model on pure roads to mixed-autonomy settings as shown in Fig. 2 (a). Assuming that neither the nominal velocity nor the maximum density changes, the critical density at which traffic becomes congested will now shift and increase with *autonomy level* of the road. This can be explained as autonomous vehicles do not require a headway as large as that needed by human drivers due to platooning benefits.

To formalize the relationship between autonomy level and critical density on road  $i$ , we assume at nominal velocity autonomous vehicles and humans require headway  $\bar{h}_i$  and  $h_i$ , respectively, with  $\bar{h}_i \leq h_i$ . This inequality reflects the assumption that autonomous vehicles can maintain a short headway, regardless of the type of vehicle they are following. Then, if  $x_i$  and  $y_i$  denote the flow of human-driven and autonomous vehicles respectively, the critical density is

$$\rho_i^{\text{crit}}(x_i, y_i) = \frac{1}{\frac{y_i}{x_i + y_i} \bar{h}_i + \frac{x_i}{x_i + y_i} h_i + L},$$

<sup>5</sup> Though congested flow may be unstable, it can be stabilized with a small number of autonomous vehicles [7], adding an additional constraint to later optimizations.

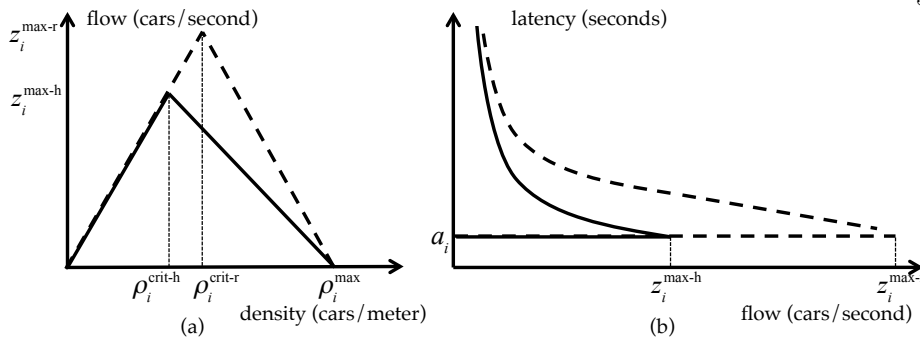


Fig. 2: **(a)** The fundamental diagram of traffic for roads with all human-driven (solid) and all autonomous (dashed) vehicles. Congestion begins at a higher vehicle density as autonomous vehicles require a shorter headway when following other vehicles. Superscripts -h and -r denote parameters corresponding to purely human and purely autonomous traffic, respectively. **(b)** The relationship between vehicle flow and latency also changes in the presence of autonomous vehicles. Free-flow speed remains the same but maximum flow on a road increases.

where all vehicles have length  $L$ . Here the denominator represents the average length from one car's rear bumper to the preceding car's rear bumper when all cars follow the vehicle in front of them with nominal headway. Note that critical density is expressed here as a function of flows  $x_i$  and  $y_i$ . This quantity can also be expressed as a function of *autonomy level*  $\alpha_i(x_i, y_i) = \frac{y_i}{x_i + y_i}$ , i.e., the fraction of vehicles on road  $i$  that are autonomous. Since flow increases linearly with density until hitting  $\rho_i^{\text{crit}}$ , the maximum flow can also be expressed as a function of autonomy level:  $z_i^{\text{max}}(x_i, y_i) = v_i^f \rho_i^{\text{crit}}(x_i, y_i)$ .

The latency on a road is a function of vehicle flows on the road as well as a binary argument  $s_i$ , which indicates whether the road is congested:

$$\ell_i(x_i, y_i, s_i) = \begin{cases} \frac{d_i}{v_i^f}, & s_i = 0, \\ d_i \left( \frac{\rho_i^{\text{max}}}{x_i + y_i} + \frac{\rho_i^{\text{crit}}(x_i, y_i) - \rho_i^{\text{max}}}{z_i^{\text{max}}(x_i, y_i)} \right), & s_i = 1. \end{cases} \quad (1)$$

We define the free-flow latency on road  $i$  as  $a_i := d_i/v_i^f$ , where  $d_i$  is the length of road  $i$ . Fig. 2 (b) illustrates the effect of mixed autonomy on latency.

**Demonstrative Example.** We use the following running example to intuitively describe the concepts formalized in Main Results (Sec. 3). Imagine driving on a Friday afternoon in Los Angeles, where you plan to drive from the Beverly Hills library to the Valley. The most direct route, Coldwater Canyon, takes 25 minutes without traffic. Taking the 405 freeway would be 30 minutes, and Laurel Canyon would be 35 minutes in free-flow. Unfortunately, it is Friday afternoon and the flow is anything but free! Let us assume that everyone is in the same predicament – meaning that all traffic on these roads is from people with the same start and end points. People would




Fig. 3: The map with three possible routes for the example.

only take Laurel Canyon if Coldwater and the 405 were congested to the point that each take at least 35 minutes. Further, any route that people use will have the same latency, otherwise they would switch to the quicker route. Fig. 4 illustrates three such equilibria with varying delays.




Fig. 4: Illustration of equilibria in the network. Equilibria can have one road in free-flow and others congested (a) or all used roads can be congested (b, c).

Krichene et al. have shown that for any given volume of traffic composed of only human drivers, there exists one *best equilibrium* which is an equilibrium where one road is in free-flow and all other used roads are congested and the number of roads used is minimized [30]. *However, when there are multiple types of vehicles that affect congestion differently, as we see in mixed-autonomy, there can be multiple equilibria with the same aggregate delay.*

Concretely, autonomous vehicles can form *platoons* to increase the free-flow capacity of a road, an effect that is accentuated on freeways. This effect can be used not only to find equilibria that minimize aggregate delay, but also to find equilibria, within the set of delay-minimizing equilibria, that can accommodate extra unforeseen flow demand. As an example, Fig. 5 (a) and (b) have the same delay, whether all autonomous flow (green) is the 405 or is split between the 405 and Coldwater. However, (b) has higher autonomy level on the 405 than (a) and can therefore accommodate more additional flow on the 405, since it matches its autonomous vehicles with the roads that benefit most from them. If a social planner can dictate equilibrium routing but does not have perfect information about flow demands, using a routing such as in (b) can make the routing more *robust* to unforeseen demand.

The social cost can further be improved if users of autonomous vehicles are *altruistic* and are willing to endure a delay longer than the quickest route available in order to improve traffic. Assume the autonomous cars are owned by entities concerned for the public good and consider a population of autonomous car users who, in exchange for cheaper rides, are willing to accept a delay of 20% longer than the quickest available route. In this case, a social planner can send some autonomous users on the 405 so that Coldwater is in an uncongested state, yielding a social cost lower than that of the best-case selfish equilibria (Fig. 5 (c)).

Before presenting our main results, we also define our network and objective. **Network Objectives.** We consider a network of  $N$  parallel roads, with indices  $\{1, \dots, n\}$ . For convenience we assume no two roads have identical free-flow latencies, and accordingly, order our roads in order of increasing free-flow latency. We use  $[m] = \{1, \dots, m\}$  to denote the set of roads with indices 1 through  $m$ .




Fig. 5: (a) and (b) represent two selfish equilibria with the same social cost. The equilibrium with all autonomous flow on the 405 freeway (b) is more robust to additional unforeseen flow demand. Solid (resp. dashed) lines characterize the road with only human drivers (resp. autonomous vehicles). The 405 has maximum flow that benefits more from autonomy than the canyons – by routing all autonomous traffic onto the 405, a social planner can make the routing more robust to unforeseen demand. If some autonomous users are altruistic, a social planner can send them on the 405 while Coldwater is uncongested (c), leading to a lower overall travel time.

We consider an inelastic demand of  $\bar{x}$  and  $\bar{y}$  (human-driven and autonomous vehicle flow, respectively), meaning that the total flow of vehicles wishing to travel across the network is independent of the delays on the roads. We describe the network state by  $(\mathbf{x}, \mathbf{y}, \mathbf{s})$ , where  $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\geq 0}^N$  and  $\mathbf{s} \in \{0, 1\}^N$ . A *feasible routing* is one for which  $x_i + y_i \leq z_i^{\max}(x_i, y_i)$  for all roads, where  $z_i^{\max}$  denotes the maximum flow on road  $i$ . We denote total flow on a road by  $z_i := x_i + y_i$ . We are interested in finding a routing, i.e. allocation of vehicles into the roads, that minimizes the total latency experienced by all vehicles,  $C(\mathbf{x}, \mathbf{y}, \mathbf{s}) = \sum_{i \in [n]} (x_i + y_i) \ell_i(x_i, y_i, s_i)$ , while satisfying the demand, i.e.  $\sum_{i \in [n]} x_i = \bar{x}$  and  $\sum_{i \in [n]} y_i = \bar{y}$ . Further, we constrain this optimization based on *selfishness* or *altruism* of the vehicles, which we will define in Section 3.

### 3 Main Results

We now make precise the aforementioned notions of selfishness, robustness, and altruism. We develop properties of the resulting equilibria, and using those, provide polynomial-time algorithms for computing optimal vehicle flows.

**Selfishness.** Human drivers are often perceived as selfish, i.e., they will not take a route with long delay if a quicker route is available to them. If all drivers are selfish this leads to a *Nash*, or *Wardrop Equilibrium*, in which no driver can achieve a lower travel time by unilaterally switching routes [17, 18]. In the case of parallel roads, this means that all selfish users experience the same travel time.

**Definition 1.** *The longest equilibrium road is the road with maximum free-flow latency that has delay equal to the delay experienced by selfish users. Let  $m_{\text{EQ}}$  denote the index of this road. We then use  $NE(\bar{x}, \bar{y}, m_{\text{EQ}})$  to denote the set of Nash Equilibria with longest equilibrium road having index  $m_{\text{EQ}}$ .*

**Definition 2.** The **longest used road** is the road with maximum free-flow latency that has positive vehicle flow of any type on it. We use  $m_{\text{ALL}}$  to denote the index of this road; if all vehicles in a network are selfish then  $m_{\text{EQ}} = m_{\text{ALL}}$ .

We define the set of **Best-case Nash Equilibria (BNE)** as the set of feasible routings in equilibrium that minimize the total latency for flow demand  $(\bar{x}, \bar{y})$ , denoted  $\text{BNE}(\bar{x}, \bar{y})$ . The following theorem provides properties of the set of BNE for mixed-autonomy roads (for pure roads see [30]).

**Theorem 1.** There exists a road index  $m_{\text{EQ}}^*$  such that all routings in the set of BNE have the below properties. Further, this index  $m_{\text{EQ}}^*$  is the minimum index such that a feasible routing can satisfy the properties:

1. road  $m_{\text{EQ}}^*$  is in free-flow,
2. roads with index less than  $m_{\text{EQ}}^*$  are congested with latency  $a_{m_{\text{EQ}}^*}$ , and
3. all roads with index greater than  $m_{\text{EQ}}^*$  have zero flow.

The running example in Fig. 5 shows that there does not necessarily exist a unique BNE. However, as estimating flow demand can be difficult, we prefer to choose a BNE that can best incorporate additional unforeseen demand. We thus develop the notion of *robustness*.

**Definition 3.** The robustness of a routing in the set of Nash Equilibria is how much additional traffic (as a multiple of the original demand flow), at the overall autonomy level, can be routed onto the free-flow road. Formally, if  $(\mathbf{x}, \mathbf{y}) \in \text{NE}(\bar{x}, \bar{y}, m)$ , then it has robustness:

$$\beta(\mathbf{x}, \mathbf{y}, \bar{x}, \bar{y}, s_m) := \begin{cases} \max \gamma \text{ s.t. } x_m + y_m + \gamma(\bar{x} + \bar{y}) \leq z_m^{\text{max}}(x_m + \gamma\bar{x}, y_m + \gamma\bar{y}) & s_m = 0 \\ 0 & s_m = 1. \end{cases} \quad (2)$$

We use **Robust Best Nash Equilibria (RBNE)** to refer to the subset of BNE that maximize robustness. The RBNE may not be unique.

**Altruism.** It is possible that some passengers, especially those in autonomous vehicles, can be incentivized to use routes that are not quickest for that individual, but instead lead to a lower social cost (Fig. 5). We use the term *altruism profile* to refer to the distribution of the degree to which autonomous vehicles are willing to endure longer routes. For computational reasons, we consider altruism profiles with a finite number of altruism levels.




Fig. 6: Altruism profiles. A fraction  $\varphi(\kappa)$  of autonomous users will not accept delay greater than  $\kappa$  times that of the quickest available route. (a) Users will tolerate delay of up to  $\kappa_0$  times that of the quickest route. (b) Users have multiple altruism levels.



Formally, we define  $\varphi : \mathbb{R}_{\geq 0} \rightarrow [0, 1]$  to represent the altruism profile as a nondecreasing function of a delay value that is mapped to  $[0, 1]$ . A volume of  $\varphi(\kappa)\bar{y}$  autonomous flow will reject a route incurring delay  $\kappa$  times the minimum route delay available, which we denote by  $\hat{\ell}_0$ . If autonomous users have a uniform altruism level as in Fig. 6 (a), we call them  $\kappa_0$ -altruistic users, where  $\kappa_0$  is the maximum multiple of the minimum delay that autonomous users will accept. Users may have differing altruism levels, as in Fig. 6 (b). We use  $K$  to denote the set of altruism levels, with cardinality  $|K|$ . Accordingly, a feasible routing  $(\mathbf{x}, \mathbf{y}, \mathbf{s})$  is in the set of **Altruistic Nash Equilibria (ANE)** if

1. all routes with human traffic have latency  $\hat{\ell}_0 \leq \ell_i(x_i, y_i, s_i) \forall i \in [n]$  and
2. for any  $\ell \geq 0$ , a volume of at least  $\varphi(\ell/\hat{\ell}_0)\bar{y}$  autonomous traffic experiences a delay less than or equal to  $\ell$ . Note that it is enough to check this condition for  $\ell = \ell_i(x_i, y_i, s_i)$  for all  $i$ .<sup>6</sup>

We denote the set of routings at Altruistic Nash Equilibria with demand  $(\bar{x}, \bar{y})$ , equilibrium latency  $\hat{\ell}_0$ , and altruism profile  $\varphi$  as  $ANE(\bar{x}, \bar{y}, \hat{\ell}_0, \varphi)$ . The set of **Best-case Altruistic Nash Equilibria (BANE)** is the subset of ANE with routings that minimize total latency.

**Theorem 2.** *There exist a longest equilibrium road  $m_{\text{EQ}}^*$  and a longest used road  $m_{\text{ALL}}^*$  with  $m_{\text{EQ}}^* \leq m_{\text{ALL}}^*$ , such that all routings in BANE have the following properties:*

1. roads with index less than  $m_{\text{EQ}}^*$  are congested,
2. roads with index greater than  $m_{\text{EQ}}^*$  are in free-flow,
3. roads with index greater than  $m_{\text{EQ}}^*$  and less than  $m_{\text{ALL}}^*$  have maximum flow.

*Remark 1.* One may be tempted to think that, like in BNE, road  $m_{\text{EQ}}^*$  will be in free-flow. Further, one might think that  $m_{\text{EQ}}^*$  is the minimum index such that all selfish traffic can be feasibly routed at Nash Equilibrium. We provide counterexamples to these conjectures in this paper's extended version [14].

**Finding the Best-case Nash Equilibria.** In general, the Nash Equilibrium constraint is a difficult combinatorial constraint. Theorem 1 however states that we can characterize the congestion profile of the roads by finding the minimum free-flow road such that Nash Equilibrium can be feasibly achieved. Once this is found, we select a feasible routing that maximizes robustness. This is formalized as follows: find the minimum  $m_{\text{EQ}}$  such that  $NE(\bar{x}, \bar{y}, m_{\text{EQ}})$  is nonempty. Then the RBNE is the set of routings that maximize the robustness of road  $m_{\text{EQ}}$ :

$$m_{\text{EQ}}^* = \arg \min_{m_{\text{EQ}} \in [n]} a_{m_{\text{EQ}}} \text{ s.t. } NE(\bar{x}, \bar{y}, m_{\text{EQ}}) \neq \emptyset, \text{ then}$$

$$BNE(\bar{x}, \bar{y}, m_{\text{EQ}}^*) = \arg \max_{(\mathbf{x}, \mathbf{y}, \mathbf{s}) \in NE(\bar{x}, \bar{y}, m_{\text{EQ}}^*)} \beta(\mathbf{x}, \mathbf{y}, \bar{x}, \bar{y}, \mathbf{s}, m_{\text{EQ}}^*). \quad (3)$$

**Theorem 3.** *Finding a solution to (3) is equivalent to finding a routing in the set of RBNE, if any exist. Further, (3) can be solved in  $O(N^4)$  time.*

<sup>6</sup> This is similar to the notion of *epsilon-approximate Nash Equilibrium* [31], but with populations playing strategies that bring them within some factor of the best strategy available to them, with each population having a different factor.

We provide a sketch of the algorithm for finding an element of the RBNE:

1. Search over  $m_{\text{EQ}}$ , the longest equilibrium road, to find the minimum  $m_{\text{EQ}}$  that is feasible for required flow demands at Nash Equilibrium.
2. For this  $m_{\text{EQ}}$ , find the routing that maximizes the robustness of road  $m_{\text{EQ}}$ . This can be formulated as a linear program, which has cubic computational complexity in the number of roads [32].

The following optimization (which can be converted to a linear program) maximizes robustness while at feasible free-flow equilibrium.

$$\begin{aligned} \max_{\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\geq 0}^N, \gamma \geq 0} \gamma \text{ s.t. } & \sum_{i \in [m_{\text{EQ}}^*]} x_i = \bar{x}, \quad \sum_{i \in [m_{\text{EQ}}^*]} y_i = \bar{y}, \quad \ell_i(x_i, y_i, 1) = a_{m_{\text{EQ}}^*} \quad \forall i \in [m_{\text{EQ}}^* - 1], \\ & x_{m_{\text{EQ}}^*} + y_{m_{\text{EQ}}^*} + \gamma(\bar{x} + \bar{y}) \leq z_{m_{\text{EQ}}^*}^{\max}(x_{m_{\text{EQ}}^*} + \gamma\bar{x}, y_{m_{\text{EQ}}^*} + \gamma\bar{y}). \end{aligned}$$

**Finding the Best-case Altruistic Nash Equilibria.** To find an element of the BANE, we need to solve:

$$\arg \min_{m_{\text{EQ}} \in [n], \hat{\ell}_0 \in [a_{m_{\text{EQ}}}, a_{m_{\text{EQ}}+1}], (\mathbf{x}, \mathbf{y}, \mathbf{s}) \in \text{ANE}(\bar{x}, \bar{y}, \hat{\ell}_0, \varphi)} C(\mathbf{x}, \mathbf{y}). \quad (4)$$

As demonstrated in [14], the BANE no longer has monotonicity in  $m_{\text{EQ}}$ , the index of the longest equilibrium road. Further, road  $m_{\text{EQ}}$  may not be in free-flow in the set of BANE. We do however have a degree of monotonicity in the latency on road  $m_{\text{EQ}}$ : for a fixed  $m_{\text{EQ}}$ , optimal social cost is minimized at the minimum feasible latency on road  $m_{\text{EQ}}$ .

**Theorem 4.** *Finding a solution to (4) is equivalent to finding a routing in the set of BANE, if any exist. Further, (4) can be solved in  $O(|K|N^5)$  time, where  $|K|$  is the number of altruism levels of autonomous vehicle users.*

In the full proof we show that for any given  $m_{\text{EQ}}$ , there are a maximum of  $|K|N$  critical points to check for the equilibrium latency,  $\hat{\ell}_0$ . We sketch the algorithm for finding a routing in the set of BANE:

1. Enumerate through all options for  $m_{\text{EQ}}$ .
2. For each  $m_{\text{EQ}}$ , enumerate through the  $|K|N$  critical points of  $\hat{\ell}_0$ .
3. For each of these combinations, find the routing that maximizes the autonomous flow on roads  $[m_{\text{EQ}}]$ , while incorporating all human flow on these roads. This can be formulated as a linear program. Then find optimal autonomous routing on the remaining roads (order  $N$  calculations).

**Improvement can be unbounded.** We motivate the schemes described above by demonstrating that without them, aggregate latency can be arbitrarily worse than with them (Fig. 7). First we show the cost at BNE can be arbitrarily worse than that at BANE, even when autonomous users have arbitrarily low altruism.

Consider three roads with free-flow latencies  $a_3 > a_2 > a_1$ , and human-driven and autonomous flow demands  $\bar{x}$  and  $\bar{y}$  relative to maximum flows  $z_1^{\max}(\bar{x}, 0) = \bar{x}$ ,  $z_2^{\max}(0, \bar{y}) = \bar{y}$ , and  $z_3^{\max}(\bar{x}, \bar{y}) \geq \bar{x} + \bar{y}$ . This means human-driven vehicles can fit on road 1 with autonomous vehicles on road 2, or both vehicle types can fit on roads 1, 2, and 3 if all are selfish, as all vehicles cannot fit on roads 1 and 2 at the same latency. Let autonomous vehicles have uniform altruism level  $\kappa_0 = \frac{a_2}{a_1}$ ,

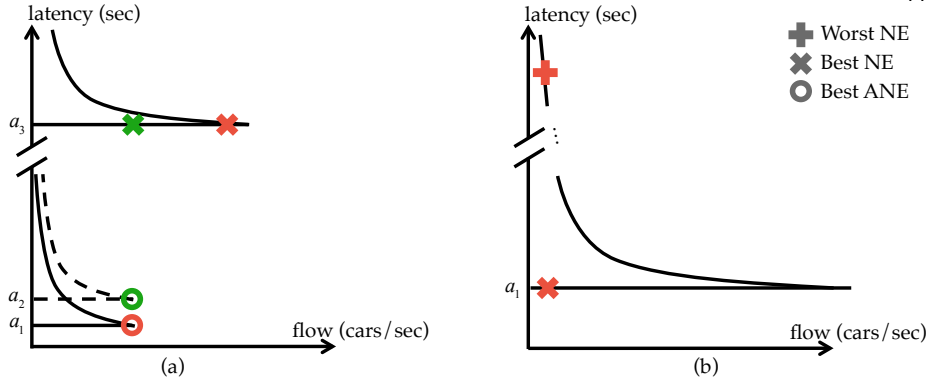


Fig. 7: Illustration showing unbounded improvement in (a) altruistic NE vs. best selfish NE, and (b) best selfish NE vs. worst NE. Red and green represent regular and autonomous vehicle flow, +’s are worst-case Nash Equilibrium routing, x’s are best-case NE routing, and circles are best-case altruistic NE routings.

which can be arbitrarily close to 1. In BANE, autonomous vehicles will be on road 2 and human-driven vehicles on road 1. However in BNE, all flow will be on road 3, as the others are fully congested. Then,  $\frac{C^{BNE}}{C^{BANE}} = \frac{a_3(\bar{x}+\bar{y})}{a_1\bar{x}+a_2\bar{y}} \geq \frac{a_3}{a_2}$ , which can be arbitrarily large.

Now consider one road with very large critical density and a small demand  $\bar{x}$ . This road can exist at free-flow latency at low density or in a highly congested state at high density and serve flow  $\bar{x}$  either way. Then,  $\lim_{\bar{x} \rightarrow 0^+} \frac{\ell_i(\bar{x}, 0, 1)}{\ell_i(\bar{x}, 0, 0)} \rightarrow \infty$ , which shows a NE can be arbitrarily worse than a BNE.

## 4 Experiments

### Numerical example of gained utility by BNE and BANE.

Next, in order to numerically show the utility gained by RBNE (or BNE) against another NE and by altruistic autonomy in a more realistic setting, we create a scenario with two parallel residential roads of length  $400\pi$  and  $1000\pi$  meters. The speed limit is 13.9 m/s ( $\sim 50$  kph). All vehicles are 5 meters long and have headway of  $\max(2, \tau v)$  meters with the vehicle in front, where  $v$  represents the speed of the vehicle and  $\tau$  is the reaction time (2 seconds for regular and 1 second for autonomous vehicles). We want to allocate the flows of 0.3 regular and 0.3 autonomous vehicles per second.

We simulated this scenario using the traffic simulator SUMO [33]. To conveniently simulate the congestion effect, we designed circular and single-lane roads (Fig. 8). We employed the original Krauss car following model for driver behavior [34]. We solved the optimizations using the CVX framework [35,36]. For comparison, we also computed and simulated a congested NE.




Fig. 8: The simulation visualization for Road 1. Regular vehicles (red) have larger headway than autonomous ones (green).

Table 1: Effect of Altruistic Autonomy

|   | NE             | RBNE         | BANE ( $\kappa_0 = 2.5$ ) |
|---|----------------|--------------|---------------------------|
| In both calculations and simulations, <b>the BNE approximately halves the total cost compared to congested NE, and the BANE halves it again</b> (Table 1). The altruism level is $\kappa_0 = 2.5$ . Gray cells represent congested roads. Subscripts $T$ and $S$ denote if the presented result is theoretical or simulated, respectively. Fig. 9 visualizes the solutions over flow-latency graph. |                |              |                           |
| $(x_1, y_1)$  | (0.006, 0.252) | (0.3, 0.031) | (0.3, 0.215)              |
| $(x_2, y_2)$  | (0.294, 0.048) | (0, 0.269)   | (0, 0.085)                |
| $C_T(\mathbf{x}, \mathbf{y})$   | 324            | 135.608      | 61.5                      |
| $C_S(\mathbf{x}, \mathbf{y})$   | 297.052        | 135.87       | 68.64                     |

halves it again (Table 1). The altruism level is  $\kappa_0 = 2.5$ . Gray cells represent congested roads. Subscripts  $T$  and  $S$  denote if the presented result is theoretical or simulated, respectively. Fig. 9 visualizes the solutions over flow-latency graph.

**Simulation of a 4-parallel-road network.** To show the properties of all the equilibria we defined, we simulated the following comprehensive scenario: We assume there are 4 parallel roads with different lengths from one point to another, two of which are residential roads where the speed limit is 13.9 m/s, and the other two are undivided highways where the limit is 25.0 m/s ( $\sim 90$  kph). The residential roads and highways have lengths  $400\pi$ ,  $600\pi$ , and  $800\pi$ ,  $1000\pi$  meters, respectively. Sorting based on the free-flow latencies makes the order of lengths  $(400\pi, 800\pi, 1000\pi, 600\pi)$ , and we use this order from this point on. We keep all other parameters the same as the previous simulation.

With this configuration, we first simulate Road 4 (residential road,  $600\pi$  meters long) alone. The resulting fundamental diagram and the relationship between vehicle flow and latency are shown in Fig. 10 along with the theoretical curves. There exists a small mismatch between the theory and simulations, which is apparent only for  $z^{\max}$ . We conjecture that this is stemming from: **1) Shape:** We approximated circular roads using pentacontagons (regular polygons with 50 edges), so the actual road lengths are slightly different. **2) Initiation:** We cannot initiate the vehicles homogeneously over the polygon. Instead, we initiate from the corners. This hinders our ability to allocate the maximum possible number of vehicles. **3) Discretization:** As the number of vehicles needs to be an integer, density values are heavily discretized. **4) Randomness:** Given the total number of vehicles for each road, our simulation initiates regular and autonomous vehicles drawn from a Bernoulli process with the given probability distribution.

To show the effect of changing  $\kappa_0$  in the altruistic case, we consider human flow demand of 0.4 and autonomous flow demand of 1.2 vehicles per second. For BANE, we simulated both  $\kappa_0 = 1.25$  and  $\kappa_0 = 1.5$ , where we assume all autonomous vehicles are uniformly altruistic for simplicity.

We then simulate the computed flow values using SUMO. We found that often the full calculated flow cannot be allocated in the simulation because of the




Fig. 9: Three solutions to the scenario with two residential roads are visualized. Blue lines represent the shorter road, whereas brown is the longer one. Solid and dashed lines correspond to fully regular and fully autonomous configurations, respectively.




Fig. 10: Simulation results of Road 4 for three different  $\alpha$  values with theoretical values (black curves). The noisy appearance of the curves for  $\alpha = 0.5$  is due to randomness, while mismatches on  $\rho^{\text{crit}}$  are due to the other imperfections. **(a)** The fundamental diagram of traffic. Due to simulation incapacities regarding initiation, we cannot simulate heavily congested regions. We used more lightly congested traffic measurements to project heavily congested cases (dashed lines). **(b)** The flow and latency relationship.

mismatch between the theoretical maximum flow on a road and the simulated maximum flow (Fig. 10 (b)). In these cases, we kept  $\alpha$ -value for that road fixed and allocated flow with respect to Nash equilibrium. Hence, for BNE, RBNE and BANE, we added the remaining flow into road  $m_{\text{EQ}}^*$ , without breaking Nash equilibrium. For NE, we again kept  $\alpha$  and the number of utilized roads fixed and swept latency values to satisfy flow conservation. These simulation results, as well as theoretically computed values for  $\mathbf{x}$ ,  $\mathbf{y}$ ,  $\beta$ ,  $C(\mathbf{x}, \mathbf{y})$  are given in Table 2.

Table 2: Theoretical and Simulated Results

|                               | NE             | BNE           | RBNE           | BANE ( $\kappa_0 = 1.25$ ) | BANE ( $\kappa_0 = 1.5$ ) |
|-------------------------------|----------------|---------------|----------------|----------------------------|---------------------------|
| $(x_1, y_1)$                  | (0.036, 0.277) | (0.075, 0.52) | (0.391, 0)     | (0.4, 0.024)               | (0.4, 0.041)              |
| $(x_2, y_2)$                  | (0.121, 0.311) | (0.2, 0.43)   | (0.009, 0.772) | (0, 0.833)                 | (0, 0.833)                |
| $(x_3, y_3)$                  | (0.161, 0.303) | (0.126, 0.25) | (0, 0.428)     | (0, 0.343)                 | (0, 0.325)                |
| $(x_4, y_4)$                  | (0.083, 0.309) | (0, 0)        | (0, 0)         | (0, 0)                     | (0, 0)                    |
| $C_T(\mathbf{x}, \mathbf{y})$ | 640            | 201.062       | 201.062        | 169.469                    | 164.56                    |
| $C_S(\mathbf{x}, \mathbf{y})$ | 599.072        | 199.985       | 199.575        | 172                        | 167.035                   |
| $\beta_T$                     | 0              | 0.183         | 0.210          | 0                          | 0                         |
| $\beta_S$                     | 0              | 0.154         | 0.135          | 0                          | 0                         |

Note that BNE and RBNE lead to lower latencies, with a slight difference in the simulation, compared to the presented NE configuration. By introducing altruistic autonomous vehicles, the overall cost can be further decreased.

While theoretical values conform with the fact that RBNE maximizes the robustness with the same cost as other BNE's, the simulated robustness for the RBNE case is smaller than that of the presented BNE solution. To understand this, we first note that the RBNE solution makes the free-flow road with the lowest latency (road 3) fully autonomous to be able to allocate more flow. However, as can be seen from Fig. 10, the mismatch between theory and simulations grows larger with increasing autonomy. This causes RBNE to be unable to allocate high prospective flow. A better maximum flow model would resolve this.

Fig. 11 shows the average speeds of vehicles on the roads and visualizes the vehicle densities for each of the theoretically presented solutions. The congested roads are black and the free-flow roads are gray. This makes clear that the NE solution leads to high latency due to heavy traffic congestion. Other solu-




Fig. 11: Roads are visualized with respect to the vehicle density values obtained from each of the solutions. Red and green dots represent regular and autonomous vehicles, respectively. We do not show headways for simplicity. Black roads are congested, while gray roads are in free-flow. On the right, average vehicle speeds (m/s) are shown.

tions support the same flows with lower vehicle densities thanks to maintaining higher speeds, yielding lower latencies. Further, the RBNE makes the free-flow road fully autonomous to enable the allocation of more additional flow. Lastly, increasing the altruism level in BANE reduces the overall traffic congestion.<sup>7</sup>

### Effect of different altruism levels.

We also analyzed the effect of altruism level under different flow constraints. Fig. 12 shows the average latency experienced by each vehicle for three different altruism levels and under different total flow requirements. We solved all cases using BANE optimization. The first heat map shows non-altruistic case where the average latencies are exactly equal to one of the free-flow latencies. This is a direct result of Theorem 1. In the altruistic cases, the average latency experienced by vehicles is lower and increasing the altruism level helps decrease the latency further.




Fig. 12: Average latency (seconds) obtained from BANE optimization for  $\bar{x}, \bar{y} \in [0, 1.5]$  cars per second and  $\kappa_0 \in \{1, 1.25, 1.5\}$ . Black regions represent infeasible demand.

## 5 Discussion

**Summary.** In this work, we develop a new model to incorporate the effect of autonomy in parallel road networks based on the fundamental diagram of traffic. We further define a notion of robustness indicating the capability of a Nash Equilibrium to be resilient to additional unforeseen flow demand. We design an optimization-based polynomial-time algorithm to find a robust Nash Equilibrium that minimizes overall latency. We then define the concept of *altruistic autonomy* to model autonomous users' willingness to accept higher latencies. We provide a polynomial-time algorithm that utilizes altruism to minimize overall latency. We demonstrate the improvements our algorithms provide using simulations.

**Limitations and Future Directions.** In this work we consider parallel networks; we wish to also study more complex transportation networks that in-

<sup>7</sup> An animated version can be found at <http://youtu.be/Hy2S6zbL6Z0> with realistic numerical values for densities, headways, car lengths, speeds, etc.

clude other modes of transportation. Congestion and altruism can be modeled for each mode of transportation. Further, in this work we use the original Krauss model [34] for car following; more complex and expressive models, such as Intelligent Driver Model [37], could yield to other interesting outcomes. This work also considers only isolated equilibria – investigating how to move the network from one equilibrium to another is a valuable direction. Finally, our work is a first step in formalizing altruistic autonomy. We would like to perform more realistic case studies to understand emergent phenomenon in mixed-autonomy roads.

## Acknowledgments

Toyota Research Institute (“TRI”) provided funds to assist the authors with their research but this article solely reflects the opinions and conclusions of its authors and not TRI or any other Toyota entity. This work was also supported by NSF grant CCF-1755808 and UC Office of the President grant LFR-18-548175.

## References

1. A. Askari, D. A. Farias et al., “Effect of adaptive and cooperative adaptive cruise control on throughput of signalized arterials,” in *IEEE Intelligent Veh.*, 2017.
2. J. Lioris, R. Pedarsani et al., “Platoons of connected vehicles can double throughput in urban roads,” *Transportation Res. Part C: Emerging Tech.*, vol. 77, 2017.
3. M. Motie and K. Savla, “Throughput analysis of a horizontal traffic queue under safe car following models,” in *IEEE Conference on Decision and Control*, 2016.
4. A. Adler, D. Miculescu, and S. Karaman, “Optimal policies for platooning and ride sharing in autonomy-enabled transportation,” in *Workshop on Algorithmic Foundations of Robotics (WAFR)*, 2016.
5. D. Lazar, K. Chandrasekher, R. Pedarsani, and D. Sadigh, “Maximizing road capacity using cars that influence people,” in *Accepted, IEEE Conference on Decision and Control*, IEEE, 2018.
6. C. Wu, A. Kreidieh, E. Vinitzky, and A. M. Bayen, “Emergent behaviors in mixed-autonomy traffic,” in *Conference on Robot Learning*, pp. 398–407, 2017.
7. C. Wu, A. M. Bayen, and A. Mehta, “Stabilizing traffic with autonomous vehicles,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1–7, IEEE, 2018.
8. S. Cui, B. Seibold, R. Stern, and D. B. Work, “Stabilizing traffic flow via a single autonomous vehicle: Possibilities and limitations,” in *IEEE Intelligent Veh.*, 2017.
9. R. E. Stern *et al.*, “Dissipation of stop-and-go waves via control of autonomous vehicles: Field experiments,” *Transportation Res. Part C: Emerging Tech.*, 2018.
10. D. Merugu, B. S. Prabhakar, and N. Rama, “An incentive mechanism for decongesting the roads: A pilot program in bangalore,” in *Proc. of ACM NetEcon Workshop*, Citeseer, 2009.
11. C. Pluntke and B. Prabhakar, “Insinc: A platform for managing peak demand in public transit,” *JOURNEYS, Land Transport Authority Academy of Singapore*, 2013.
12. Y. M. Nie and Y. Yin, “Managing rush hour travel choices with tradable credit scheme,” *Transportation Research Part B: Methodological*, 2013.
13. C. F. Daganzo, “The cell transmission model: A dynamic representation of highway traffic consistent with the hydrodynamic theory,” *Transportation Res. Part B: Methodological*, 1994.

14. E. Bıyık, D. A. Lazar, R. Pedarsani, and D. Sadigh, “Altruistic autonomy: Beating congestion on shared roads,” *arXiv preprint arXiv:1810.11978*, 2018.
15. D. W. Hearn, S. Lawphongpanich, and S. Nguyen, “Convex programming formulations of the asymmetric traffic assignment problem,” *Transportation Res. Part B: Methodological*, 1984.
16. S. C. Dafermos, “The traffic assignment problem for multiclass-user transportation networks,” *Transportation science*, vol. 6, no. 1, pp. 73–87, 1972.
17. M. Florian and D. Hearn, “Network equilibrium models and algorithms,” *Handbooks in Operations Research and Management Science*, vol. 8, pp. 485–550, 1995.
18. S. Dafermos, “Traffic equilibrium and variational inequalities,” *Transportation science*, vol. 14, no. 1, pp. 42–54, 1980.
19. T. Roughgarden, “Stackelberg scheduling strategies,” *SIAM Journal on Computing*, vol. 33, no. 2, pp. 332–350, 2004.
20. V. Bonifaci, T. Harks, and G. Schäfer, “Stackelberg routing in arbitrary networks,” *Mathematics of Operations Research*, vol. 35, no. 2, pp. 330–346, 2010.
21. C. Swamy, “The effectiveness of stackelberg strategies and tolls for network congestion games,” *ACM Transactions on Algorithms (TALG)*, 2012.
22. T. Roughgarden and É. Tardos, “How bad is selfish routing?,” *Journal of the ACM (JACM)*, vol. 49, no. 2, pp. 236–259, 2002.
23. J. R. Correa, A. S. Schulz, and N. E. Stier-Moses, “A geometric approach to the price of anarchy in nonatomic congestion games,” *Games Econ. Behavior*, 2008.
24. G. Perakis, “The price of anarchy under nonlinear and asymmetric costs,” *Mathematics of Operations Research*, vol. 32, no. 3, pp. 614–628, 2007.
25. D. A. Lazar, S. Coogan, and R. Pedarsani, “Routing for traffic networks with mixed autonomy,” *arXiv preprint arXiv:1809.01283*, 2018.
26. C. F. Daganzo, “A continuum theory of traffic dynamics for freeways with special lanes,” *Transportation Res. Part B: Methodological*, 1997.
27. H. Zhang and W. Jin, “Kinematic wave traffic flow model for mixed traffic,” *Transportation Research Record: Journal of the Transportation Research Board*, 2002.
28. S. Logghe and L. Immers, “Heterogeneous traffic flow modelling with the lwr model using passenger-car equivalents,” in *10th World congress on ITS*, 2003.
29. S. Fan and D. B. Work, “A heterogeneous multiclass traffic flow model with creeping,” *SIAM Journal on Applied Mathematics*, vol. 75, no. 2, pp. 813–835, 2015.
30. W. Krichene, J. D. Reilly, S. Amin, and A. M. Bayen, “Stackelberg routing on parallel transportation networks,” *Handbook of Dynamic Game Theory*, 2017.
31. N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, *Algorithmic game theory*. Cambridge University Press, 2007.
32. C. C. Gonzaga, “Path-following methods for linear programming,” *SIAM review*, vol. 34, no. 2, pp. 167–224, 1992.
33. D. Krajzewicz, J. Erdmann, M. Behrisch, and L. Bieker, “Recent development and applications of sumo-simulation of urban mobility,” 2012.
34. S. Krauß, P. Wagner, and C. Gawron, “Metastable states in a microscopic model of traffic flow,” *Physical Review E*, vol. 55, no. 5, p. 5597, 1997.
35. M. Grant and S. Boyd, “CVX: Matlab software for disciplined convex programming, version 2.1.” <http://cvxr.com/cvx>, Mar. 2014.
36. M. Grant and S. Boyd, “Graph implementations for nonsmooth convex programs,” in *Recent Advances in Learning and Control*, 2008.
37. M. Treiber, A. Hennecke, and D. Helbing, “Congested traffic states in empirical observations and microscopic simulations,” *Physical review E*, vol. 62, no. 2, 2000.