(In)stability of departure time choice with the bottleneck model

Raphaël Lamotte
Nikolas Geroliminis

Urban Transport Systems Laboratory (LUTS) **May 2018**

STRC | 18th Swiss Transport Research Conference
Monte Verità / Ascona, May 16 – 18, 2018

Urban Transport Systems Laboratory (LUTS)

(In)stability of departure time choice with the bottleneck model

Raphaël Lamotte, Nikolas Geroliminis
Urban Transport Systems Laboratory
EPFL
Station 18 CH-1015 Lausanne
phone: +41-21-693 24 79
fax: +41-21-693 24 69
{raphael.lamotte,nikolas.geroliminis}@epfl.ch

May 2018

Abstract

This paper investigates the stability of departure time choice during the morning commute with a bottleneck model of constant capacity. It is shown that the corresponding utility function is monotonic if and only if the marginal utility rate of time at the destination is a non-increasing function of time. Thus, the morning commute equilibrium is likely to be unstable even in continuous time for a wide range of schedule preferences and rational adjustment mechanisms. This result explains the many instances of non-converging behaviors reported in the literature. It is however observed that user sensitivity to utility differences and heterogeneity in schedule preferences have a strong influence on the size of the oscillations. If the population is sufficiently heterogeneous and insensitive to utility differences, the instability may be hardly noticeable among regular seasonal variations.

Keywords
stability, departure time, monotonicity, heterogeneity, bottleneck, equilibrium

1 Introduction

In transportation economics and transportation engineering, the "morning commute problem" refers to the study of departure time choice for a population of commuters. Commuters experience a generalized cost which depends both on the time of the day and on the duration of their trip, which depends itself on the decisions of other users. Most studies so far have focused on existence, uniqueness or on other static properties of the Nash equilibrium and of the social optimum (see Small (2015) for a review).

Despite its practical importance, the question of equilibrium stability has received very little attention. Studying stability is more challenging as it involves an additional component, referred to as the "adjustment mechanism" in the route choice literature, or as the "revision protocol" in evolutionary game theory. In simple words, the equilibrium stability depends on how users react to perturbations.

To the best of our knowledge, the stability of the morning commute equilibrium was only investigated by few authors that mostly reported negative results. Observations of non-converging behavior in simulations were reported by de Palma (2000), Mc Breen *et al.* (2006), Iryo (2008), Bressan *et al.* (2012), Guo *et al.* (2018). The analytical work is more limited. Iryo (2008) explained that some difference with the route choice problem precludes the transposition of Smith's stability proof (Smith, 1984b) to the departure time choice problem . More recently, Guo *et al.* (2018) proved discrete time instability provided that the population size is large enough. Yet, all these results only apply to specific examples of adjustment mechanisms and, except for Bressan *et al.* (2012), only to the conventional $\alpha - \beta - \gamma$ schedule preferences. The only "positive" stability results are those based on the concept of "bounded rationality", which considers that users do not update their decisions if the cost reduction that would obtained by doing so is smaller than some threshold. Yet, bounded rationality leads to a continuum of equilibria, which is often impractical for policy evaluation.

This paper extends the works previously mentioned by formulating results that are valid for a broad range of schedule preferences and adjustment mechanisms. These results build mostly on the concept of monotonic utility function, which is central to many stability results in route choice and evolutionary game theory. In line with previous works, the results obtained point to the instability of the morning commute equilibrium. We complement these results by a more quantitative analysis of discrete time stability and sensitivity to heterogeneity in schedule preferences.

Note however that the assumptions used vary throughout the paper, as we rely on results from various fields. The choice set for instance is considered to be discrete in evolutionary game theory and route choice, but continuous in the departure time choice literature. Similarly, evolutionary game theory is mostly concerned with continuous time systems and only allows inter-group heterogeneity. In the real world however, departure time decisions for the morning commute are taken at most once a day, and schedule preferences vary continuously among users. We did not attempt to transpose all theoretical results in a single framework of assumptions, because we believe most of the simplicity and elegance of these results would be lost in the process. While we do believe that the insights obtained would still hold under more realistic assumptions, readers should be aware of this limitation.

2 Stability and monotonicity: methodological framework

2.1 Problem decomposition

To fully describe a morning commute problem, one should specify the characteristics of the population considered and three functions: a congestion mechanism, schedule preferences and an adjustment mechanism. The congestion mechanism maps the departure time decisions to the corresponding arrival times (here, we consider a simple uni-directional bottleneck). The schedule preferences translate pairs of departure and arrival times into utilities. The adjustment mechanism, also known as revision protocol, defines how users update their decisions depending on the situation. By applying these three functions iteratively as illustrated in Fig. 1, one can model the evolution of departure time decisions from day to day.

This paper draws on evolutionary game theory to shed light on the stability of this dynamic process. In game theory, the game is a function that translates the user decisions into payoffs (payoff is the term used in game theory, utility is the term used in choice modeling). Thus, the game considered here is the combination of the congestion mechanism and schedule preferences.

2.2 Notations and definitions

We consider a single homogeneous population, i.e. a continuum of identical users. The demand is assumed inelastic, so that all agents depart everyday. The set of possible departure times is finite and is denoted $\mathrm{DT} = \{t_1, ...t_n\}$. The set of possible populations states is $X =$

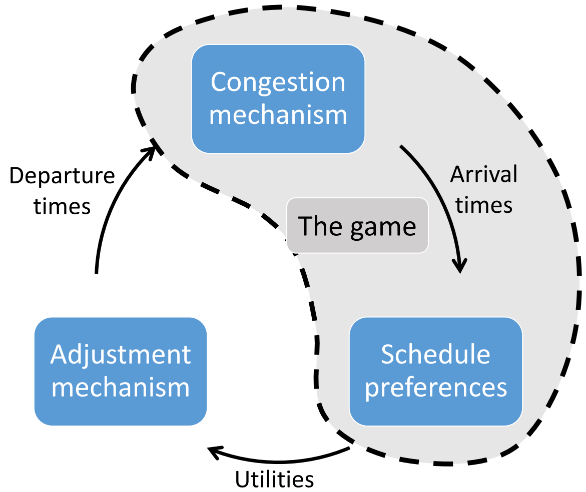

Figure 1: Ingredients of the problem

$\{(x_1, ...x_n) \in \mathbb{R}^n_+, \sum_{i=1}^n x_i = 1\}$, where x_i denotes the proportion of users choosing departure time t_i. A population game is characterized by a continuous utility function $U : X \to \mathbb{R}^n$, where $U_i(x)$ represents the utility of departing at time i when the population plays according to x. Since agents have no mass, a population state x^* is a Nash equilibrium if $x^*_i > 0$ implies that $U_i(x^*) \geq U_j(x^*)$ for all $j \in \{1, ..., n\}$ or, equivalently, if $\langle y - x^*, U(x^*)\rangle \leq 0$ for all $y \in X$.

All population games have at least one Nash equilibrium (this result was actually shown for a broader class of noncooperative games with continuums of agents, cf. Sandholm (2010b)). The concept of Nash equilibrium is also known as "user equilibrium" in the transportation literature, or as "Wardrop equilibrium" in the route choice literature.

2.3 Adjustment mechanisms

Agents receive sporadically the opportunity to update their decision and do so by following an adjustment mechanism. We focus here on reactive protocols of the form $\rho : X \times \mathbb{R}^n \to \mathbb{R}^{n\times n}_+$, which map population states $x \in X$ and their corresponding utility vectors $u \in \mathbb{R}^n$ to matrices of conditional switch rates ρ_{ij}. Assuming that the revision protocol is Lipschitz continuous, it defines a "deterministic evolutionary dynamic", i.e. a map that assigns to each continuous utility function $U : X \to \mathbb{R}^n$ a system of ordinary differential equations

$$\dot{x}_i = \sum_{j=1}^n x_j \rho_{ji}(x, U(x)) - x_i \sum_{j=1}^n \rho_{ij}(x, U(x)). \tag{1}$$

In the transportation literature, the most common protocols assume that users only consider utilities to revise their decisions. This is in contrast with evolutionary game theory, where agents are often assumed to imitate each other. "Best response" is perhaps the first utility-based revision protocol that comes to the mind: when revising their decisions, agents simply choose one of the currently optimal strategies. However, such a revision protocol is not continuous and implicitly assumes that users know the utilities of all strategies.

Another well-established revision protocol is the proportional swap introduced by Smith (1984b). It is defined by $\rho_{ij}(x,u) = \left[u_j - u_i\right]_+$ (where $\forall y \in \mathbb{R}, [y]_+ = \max(0,y)$) and it leads to the following continuous dynamic:

$$\dot{x}_i = \sum_{j=1}^{n} x_j \left[U_i(x) - U_j(x)\right]_+ - x_i \sum_{j=1}^{n} \left[U_j(x) - U_i(x)\right]_+ . \tag{2}$$

The proportional swap can be interpreted as follows: every user regularly revises her decision by comparing her current strategy against a randomly selected alternative. If the alternative provides a larger utility, the user adopts it with a probability that is proportional to the utility difference. Otherwise, she retains her current strategy. Note that while proportional swap is the most common, all revision protocols $\rho_{ij}(x,u) = \phi_j(u_j - u_i)$ where ϕ_j are sign-preserving functions[1] induce dynamics that have similar properties (Hofbauer and Sandholm, 2009). These dynamics are referred to as "impartial pairwise comparison" (IPC) dynamics.

Although we do not use them in this paper, we shall also mention the adjustment mechanisms of the form $\rho_{ij}(x,u) = \phi_j\left(u_j - \sum_{k=1}^{n} x_k u_k\right)$ (where the functions ϕ_j are sign-preserving) according to which the users switch to (resp. abandon) all strategies that have a utility larger (resp. smaller) than the average payoff experienced by the population. The corresponding dynamics are called "separable excess payoff" (SEP) dynamics. For more details and more examples of adjustment mechanisms, the reader is referred to Sandholm (2015).

Note that with the IPC and SEP classes of dynamics, $\dot{x}$ is strongly dependent on the units considered for utility. While our numerical simulations are all based on the proportional swap dynamics, we varied the sensitivity to utility differences by introducing a multiplying factor λ, such that all utilities are multiplied by λ in Eq. (2).

[1] Note that $\rho_{ij}(x,u)$ is always non-negative, by definition of conditional switch rates. Thus, the sign-preserving constraint ensures that $u_j \leq u_i \Leftrightarrow \phi_j(u_j - u_i) = 0$.

2.4 Stability: definitions and useful results

We now introduce some basic concepts from the theory of dynamical systems.

Definition 1 (stable equilibrium)**.** *A state x_e is Lyapunov stable if for every $\epsilon > 0$, there exists $\delta > 0$ such that, if $\|x(0) - x_e\| < \delta$ then for every $t > 0$, $\|x(t) - x_e\| < \epsilon$.*

Definition 2 (asymptotically stable equilibrium)**.** *A state x_e is asymptotically stable if it is Lyapunov stable and there exists $\delta > 0$ such that if $\|x(0) - x_e\| < \delta$, then $\lim_{t\to\infty} \|x(t) - x_e\| = 0$.*

Furthermore, if a Lyapunov stable equilibrium x_e is such that $\lim_{t\to\infty} \|x(t) - x_e\| = 0$ for any $x(0)$, then it is said to be *globally asymptotically stable*.

Note that the stability of a Nash equilibrium depends both on the game and on the revision protocol. In the real world, we usually know very little about the way agents take decisions. This is in sharp contrast with mechanical engineering, where the trajectory of a solid is uniquely determined by the different forces involved. Fortunately, previous research in evolutionary game theory has identified general properties of games and revision protocols, that if satisfied, are sufficient to characterize the behavior of the whole system.

2.4.1 Monotonicity and global stability results

Definition 3 (Monotonic utility function)**.** *A utility function $U : X \to \mathbb{R}^n$ is monotonic if $\langle y - x, U(y) - U(x)\rangle \leq 0$ for all $x, y \in X$.*

In plain words, monotonicity requires that the average utility improvement of alternatives that are abandoned (weighted by how many users abandon them) is larger than the (also weighted) average utility improvement of alternatives that agents are switching to (improvements can be positive or negative). We now briefly summarize some important results related to monotonic functions. These results are due to various authors that we do not attempt to list here, but the reader can refer to Hofbauer and Sandholm (2009)) for more details.

An important consequence of monotonicity is that the set of Nash equilibria is convex. If in addition the utility function is strictly monotonic at some Nash equilibrium x^* (i.e. if $\langle y - x^*, U(y) - U(x^*)\rangle < 0$ for all $y \in X\backslash\{x\}$), then this Nash equilibrium is unique. Provided that the utility function is continuously differentiable, monotonicity also guarantees that the set of Nash equilibria is globally attracting for multiple dynamics, including all those introduced in Section 2.3. If the set of Nash equilibria is reduced to a singleton, then it is globally asymptotically stable. Note that even though the results formulated are rather general, the stability conditions stated are only sufficient, and not necessary.

Games with monotonic utility functions are known as "contractive games" or "stable games". For instance, games such that for all x and for all $i \in \{1, ...n\}$, $\frac{\partial U_i}{\partial x_i}(x) \leq 0$ and $\frac{\partial U_i}{\partial x_i}(x) \leq \sum_{j \neq i} \left| \frac{\partial U_i}{\partial x_j}(x) \right|$ (i.e. negative diagonal dominant games) are such that the matrix $DU(x)$ is negative semidefinite, and therefore are contractive games.

2.4.2 Local stability

The concept of Evolutionarily Stable State introduced by Maynard Smith and Price (1973) is a local analogue of a monotonic utility function.

Definition 4 (Evolutionarily Stable State)**.** *A state x is an Evolutionarily Stable State (ESS) of a continuous utility function U if it is a Nash equilibrium for this function U and if there is a neighborhood O of x such that for all $y \in O \backslash \{x\}$, $\langle y - x, U(x) \rangle = 0$ implies that $\langle y - x, U(y) \rangle < 0$.*

In other words, if users can be equally well-off by playing according to y when the population state is at the equilibrium x, then users would be better off playing according to x if the population state was y. Intuitively, this definition already involves some concept of stability (if the system moves from x to y, it seems likely to come back to x), even though it characterizes the utility function only, and not the revision protocol.

While some stability results can already be obtained with the concept of ESS, the main theorem used in this paper relies on the stronger concept of a regular ESS (due to Taylor and Jonker (1978)). Let us define the support of a population state $x \in X$ as $\text{supp}(x) = \{i \in \{1, ...n\}, x_i > 0$ and the tangent space to X, $TX = \{z \in \mathbb{R}^n, \sum_{i=1}^n z_i = 0\}$.

Definition 5 (regular ESS)**.** *A state x is an regular Evolutionarily Stable State (ESS) of a differentiable utility function U if*

- *it is a quasistrict Nash equilibrium, i.e. $U_i(x) = \max_{j=1,...n}(U_j(x)) \Leftrightarrow x_i > 0$;*
- *$z' DU(x) z < 0$ for all $z \in TX \backslash \{0\}$ such that $supp(z) \subset supp(x)$.*

Compared to an ESS, a regular ESS requires that all unused strategies are suboptimal for users and that the stability condition $\langle y - x, U(y) \rangle < 0$ holds after linearizing U at x. We then have the following result:

Theorem 1 (Sandholm (2010a)). *Let x^* be a regular ESS of U. Then x^* is asymptotically stable under*

1. *any impartial pairwise comparison dynamic for U;*
2. *any separable excess payoff dynamic for U;*
3. *the best response dynamic for U.*

2.5 Additional intuition on monotonicity and convergence in discrete-time

Results on discrete-time stability typically require a system of the form $x^{(k)} = Ax^{(k-1)}$ (either valid globally, or as a linearization around the equilibrium). We do not attempt to obtain such a model here as (i) reasonable adjustment mechanisms are typically not continuously differentiable around the equilibrium and (ii) we aim at understanding the importance of the congestion mechanism, schedule preferences and adjustment mechanism separately. While usage of the monotonicity property is usually restricted to the context of continuous-time stability, we explain hereafter that monotonicity can also provide a rich intuition concerning discrete-time stability issues.

Let us consider an interior equilibrium x^*, a nearby x state resulting from some perturbation, as well as the corresponding utility vectors u^* and u.[2] Let u_T^* and u_T be the images of u^* and u by the orthogonal projection on the space tangent to X, TX. Since x^* is an interior equilibrium, all alternatives must have the same utility, so u_T^* is the null vector. Given that $x - x^* \in TX$, the monotonicity condition reduces to $\langle x - x^*, u_T \rangle \leq 0$. Thus, with a monotonic utility function, if the state followed exactly the projection of the utility vector into the tangent space of feasible population changes TX, the Euclidian distance from x^* would be decreasing with time.

In discrete time however, the relation between the vectors $x - x^*$ and u_T should not be reduced to their scalar product. If the vector u_T is almost perpendicular to $x - x^*$ and the system follows approximately its direction, the system will either get further from x^*, or it will get closer but extremely slowly, depending on the user sensitivity λ.

In reality, the system's direction depends on the adjustment mechanism. Yet, the proofs of Lyapunov stability proposed by Smith (1984b), Mounce (2006), Hofbauer and Sandholm (2009) follow a similar idea. They consider a function V that is a measure of distance to equilibrium

[2] The reasoning is essentially the same when the equilibrium is a corner point, except that we should focus on the dimensions i such that either $x_i > 0$ or $x_i^* > 0$.

and that it is shown to be a Lyapunov function, i.e. that (i) $V(x) \geq 0$ for all $x \in X$, (ii) $V(x) = 0$ if and only if x is an equilibrium, and (iii) $\langle \nabla V(x), \Phi(x) \rangle < 0$ if x is not an equilibrium. While it is easy to find a function V that satisfies the requirements (i) and (ii), requirement (iii) can be more challenging as the possible choices of V differ for every adjustment mechanism. Yet, for all three families of adjustment mechanisms introduced in Section 2.3, requirement (iii) is proven by showing that

$$\langle \nabla V(x), \Phi(x) \rangle = \mu \langle \dot{x}, DU(x)\dot{x} \rangle + A(x), \tag{3}$$

where $\mu > 0$, $DU(x)$ is the jacobian matrix of U at x and $A(x)$ is a function that is specific to the adjustment mechanism chosen that is strictly negative for all x that is not an equilibrium. The monotonicity of U guarantees that $\langle \dot{x}, DU(x)\dot{x} \rangle \leq 0$, so the requirement (iii) is satisfied. For more details, the reader can refer to the proof of the stability theorem in Smith (1984b) or to the proofs of Theorem 5.3 and 7.1 in Hofbauer and Sandholm (2009).

For instance for the Proportional Swap of Smith (1984b), a possible function V is

$$V(x) = \frac{1}{2} \sum_{i=1}^{n} x_i \sum_{j=1}^{n} \left(\left[U_j(x) - U_i(x) \right]_+ \right)^2. \tag{4}$$

The corresponding function A is $A(x) = \frac{1}{2} \sum_{i=1}^{n} \dot{x}_i \sum_{j=1}^{n} \left(\left[U_j(x) - U_i(x) \right]_+ \right)^2$, where $\dot{x}$ is given by Eq. (1). It takes a few additional steps to prove that $A(x) < 0$ for all states x that are not equilibria but it is already clear that the magnitude of $A(x)$ decreases as the differences between used alternatives becomes smaller. Thus, the influence of monotonicity on the system evolution can be expected to be stronger near the equilibrium.

3 Stability of the bottleneck model

3.1 The bottleneck model

We consider a single origin-destination pair, connected by a single route with a bottleneck of constant capacity. As it is common in the literature, we ignore the constant component of travel time and consider only the queuing-related delays (without loss of generality). The delay experienced by a user that arrives at the bottleneck at time t is given by $d(t) = q(t)/s$, where $q(t)$ is the number of vehicles queuing at the bottleneck at time t and s is the bottleneck capacity (in

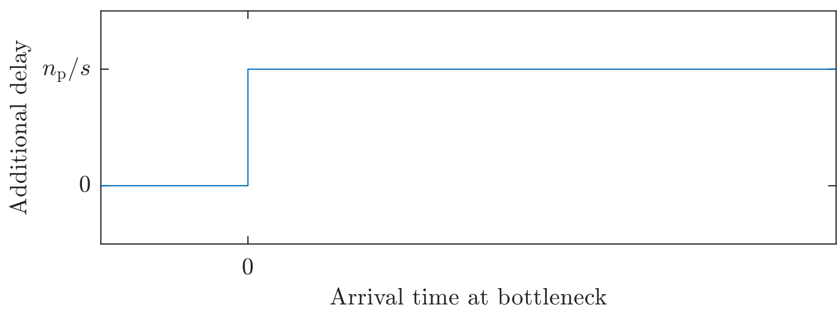

Figure 2: Additional delay imposed by a user of mass n_p arriving at time $t = 0$ at a bottleneck of capacity s on others, as a function of others' arrival time at the bottleneck. The bottleneck is assumed to be congested for the entire period of interest.

veh/h). If we denote by r the arrival rate at the bottleneck, the queue length evolution obeys

$$\dot{q}(t) = \begin{cases} r(t) - s, & \text{if } q(t) > 0 \text{ or } r(t) > s \\ 0, & \text{otherwise.} \end{cases} \tag{5}$$

The congestion mechanism described by the constant capacity bottleneck model implies highly asymmetric externalities: users are only delayed by those traveling before them, and only delay those traveling after them. Moreover, externalities do not vanish with time: all users traveling after some perturbation are delayed by the same amount, as long as the queue does not vanish. This is illustrated in Fig. 2.

3.2 The morning commute with the bottleneck model

Congestion in the morning commute arises from the fact that users have similar schedule preferences. We follow here Vickrey (1973), Tseng and Verhoef (2008), Fosgerau and de Palma (2012) and assume that these schedule preferences are expressed in a utility function of the form $U(t_1, t_3) = H(t_1) + W(t_2)$, where t_1 and t_2 are the times of departure (from home) and arrival (at work) and $H(t_d) = \int_0^{t_d} h(t_d)\,\mathrm{d}t$ and $W(t_a) = \int_{t_a}^{0} w(t_a)\,\mathrm{d}t$. We further assume that the marginal utility rates at home (h) and at work (w) are positive everywhere, piece-wise continuous, and that there exists t^* such that for all $t < t^*$, $h(t) > w(t)$ and for all $t > t^*$, $h(t) < w(t)$.

In the simulations hereafter, we use $h(t) = \alpha = 1$ and $w(t) = \alpha + \frac{\gamma-\beta}{2} + \frac{\beta+\gamma}{\pi}\tan^{-1}(w(t - t^*))$, where $w = 4$, $\beta = 0.5$ and $\gamma = 2$. These preferences represent a smooth approximation of the well-known $\alpha - \beta - \gamma$ preferences: they share the same marginaly utility rate at home and have similar utility rates at work, as illustrated in Fig. 3.

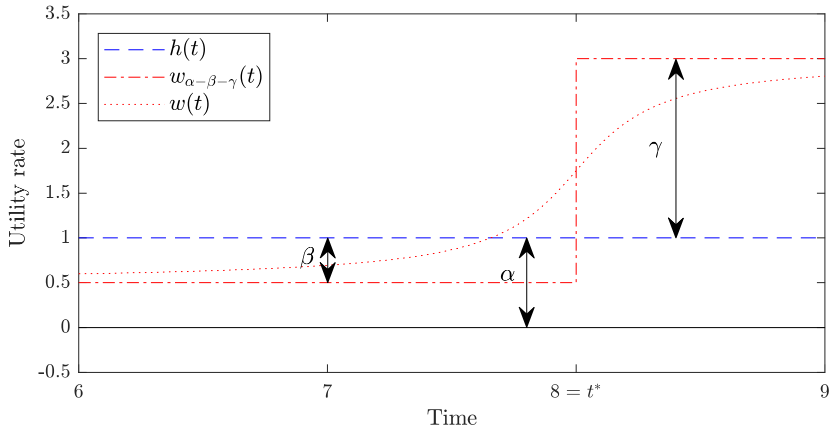

Figure 3: Schedule preferences.

Existence (resp. uniqueness) of the equilibrium has been established by Smith (1984a) (resp. Daganzo (1985)) for the case of a homogeneous population with constant h and continuous and (resp. strictly) increasing w. Various types of heterogeneity were introduced by Newell (1987) and Arnott *et al.* (1988). Lindsey (2004) extended these results to allow for more general dependence of utility on the arrival time, and Fosgerau and de Palma (2012) extended them for a general utility function depending both on departure time and arrival time, for the case of a continuously distributed free-flow travel time.

3.3 (Non)-monotonicity with homogeneous users

3.3.1 Analytical results

The monotonicity of the utility function corresponding to a constant capacity bottleneck was proven for the case without schedule preferences by Smith and Ghali (1990). This result was then extended by Mounce (2006) to handle the case of a bottleneck with time-dependent capacity.

Theorem 2 (Mounce (2006))**.** *Without schedule preferences, the bottleneck delay function is a monotonic function of the flow into the bottleneck if and only if the bottleneck capacity is non-decreasing with respect to time.*

We first provide a simple example illustrating this result and then show how a very similar proposition can be formulated with schedule preferences.

Consider a constant capacity bottleneck that is consistently congested during some time period. Consider then the modification of the departure rate illustrated in Fig. 4, occurring entirely

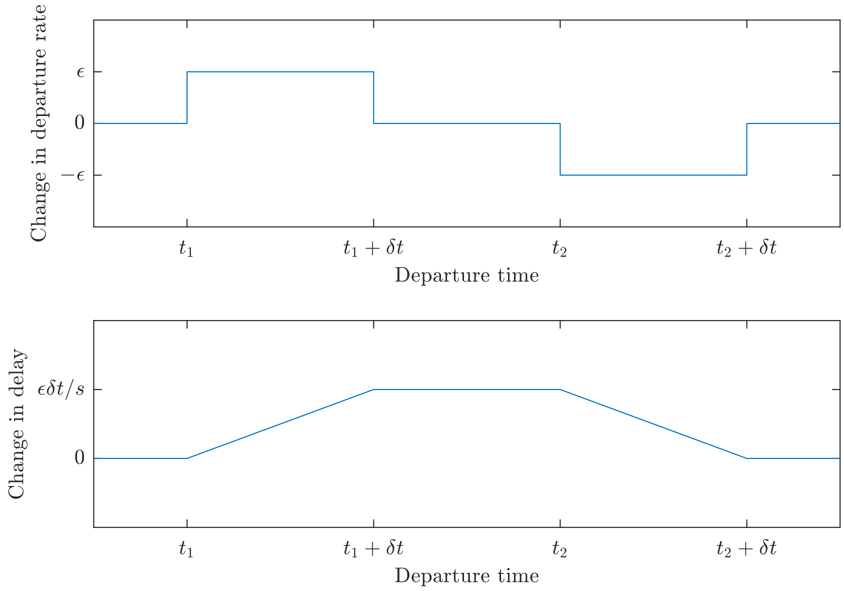

Figure 4: Influence of a small perturbation of the departure rate on experienced delays.

inside the congested period. Let r_1 and r_2 denote the original and modified departure rates. If utility is simply given by $U = -T$ (where T is the travel time), we have that $\langle r_1 - r_2, U_{r_1} - U_{r_2}\rangle = \int_{-\infty}^{\infty}(r_1(t) - r_2(t))(U_{r_1}(t) - U_{r_2}(t))\,\mathrm{d}t = \int_{t_1}^{t_1+\delta t} \frac{\epsilon^2}{s}(t - t_1)\,\mathrm{d}t + \int_{t_2}^{t_2+\delta t} -\frac{\epsilon^2}{s}(t_2 + \delta t - t)\,\mathrm{d}t = 0$. In other words, the change in utility function is perpendicular to the change in departure rate: users have exactly the same incentive to choose a strategy in the time slot $[t_1, t_1 + \delta t]$ (where users were added) and in the time slot $[t_2, t_2 + \delta t]$ (where users were removed).

In this light, it is not surprising that the global stability result of Smith and Ghali (1990) represents a limit case. If capacity slightly decreased during $[t_1, t_2 + \delta t]$, delays would increase for users arriving after the capacity decrease, leading to $\langle r_1 - r_2, U_{r_1} - U_{r_2}\rangle > 0$. Similarly, if there was even the slightest increase in marginal utility rate at the destination, the delays after this increase would be more costly for users, and the same conclusion would apply. This leads us to formulate the following corollary of Theorem 2.

Corollary 1. *Assume that $w(t) > 0$ for all t. With a bottleneck of constant capacity, the bottleneck utility function is a monotonic function of the flow into the bottleneck if and only if the marginal utility function at work w is a non-increasing function of time.*

Proof. Let $s(t)$ denote the capacity of the bottleneck at time t and let $A(t) = \int_0^t s(\tau)\,\mathrm{d}\tau$ denote the maximum possible cumulative number of arrivals at destination at time t, starting from time

0. In the general case (variable w and variable s), the utility $u(t_0)$ of departing at time t_0 is

$$u(t_0) = \begin{cases} H(t_0) + W(t_0), & \text{if } q(t_0) = 0 \\ H(t_0) + W(t_0) - \int_{t_0}^{A^{-1}(A(t_0)+q(t_0))} w(\tau)\,\mathrm{d}\tau, & \text{otherwise.} \end{cases}$$

By making the change of variable $a = A(t)$, this becomes

$$u(t_0) = \begin{cases} H(t_0) + W(t_0), & \text{if } q(t_0) = 0 \\ H(t_0) + W(t_0) - \int_{A(t_0)}^{A(t_0)+q(t_0)} \frac{w(A^{-1}(a))}{s(A^{1-}(a))}\,\mathrm{d}a, & \text{otherwise.} \end{cases}$$

Thus, the utility function does not depend on w and s separately, but only on their ratio. In particular, a problem with constant capacity s_0 and time varying marginal utility at work $w_0(t)$ has the same utility function as a problem with time varying capacity $s_0/w_0(t)$ and constant marginal utility at work equal to 1. □

Since the marginal utility rate w is expected to be an increasing function of time (Tseng and Verhoef, 2008), the utility function associated with a constant capacity bottleneck is expected not to be monotonic.

Although the focus of this paper is on the user equilibrium, note that we can also obtain a result on the stability of the social optimum at almost no additional cost. Indeed, consider the following socially optimal toll function of the passage time at the bottleneck:

$$\$(t) = \max(H(t) + W(t) - \bar{U}, 0), \tag{6}$$

where $\bar{U}$ denotes the minimum utility under a socially optimal inflow function, excluding any toll. Since the social optimum does not involve any queueing, the bottleneck should be used at capacity at all times such that $\bar{U} < H(t) + W(t)$. At such times, the utility becomes $U = H(t_d) + W(t_a) + \bar{U} - H(t_a) - W(t_a) = H(t_d) + \bar{U} - H(t_a) = H(t_d) + \bar{U} + \int_{t_a}^{0} h(t)\,\mathrm{d}t$. This utility function has the same structure as in the no-toll case, except that $h(t)$ in the tolled case plays the role of $w(t)$ in the no-toll case. We thus obtain another corollary:

Corollary 2. *With a bottleneck of constant capacity and the toll* (6)*, the bottleneck utility function is a monotonic function of the flow into the bottleneck if and only if the marginal utility function at home h is a non-increasing function of time.*

Note that unlike w, h is expected to decrease during the morning period, so that the social optimum would be globally attracting for many adjustment mechanisms.

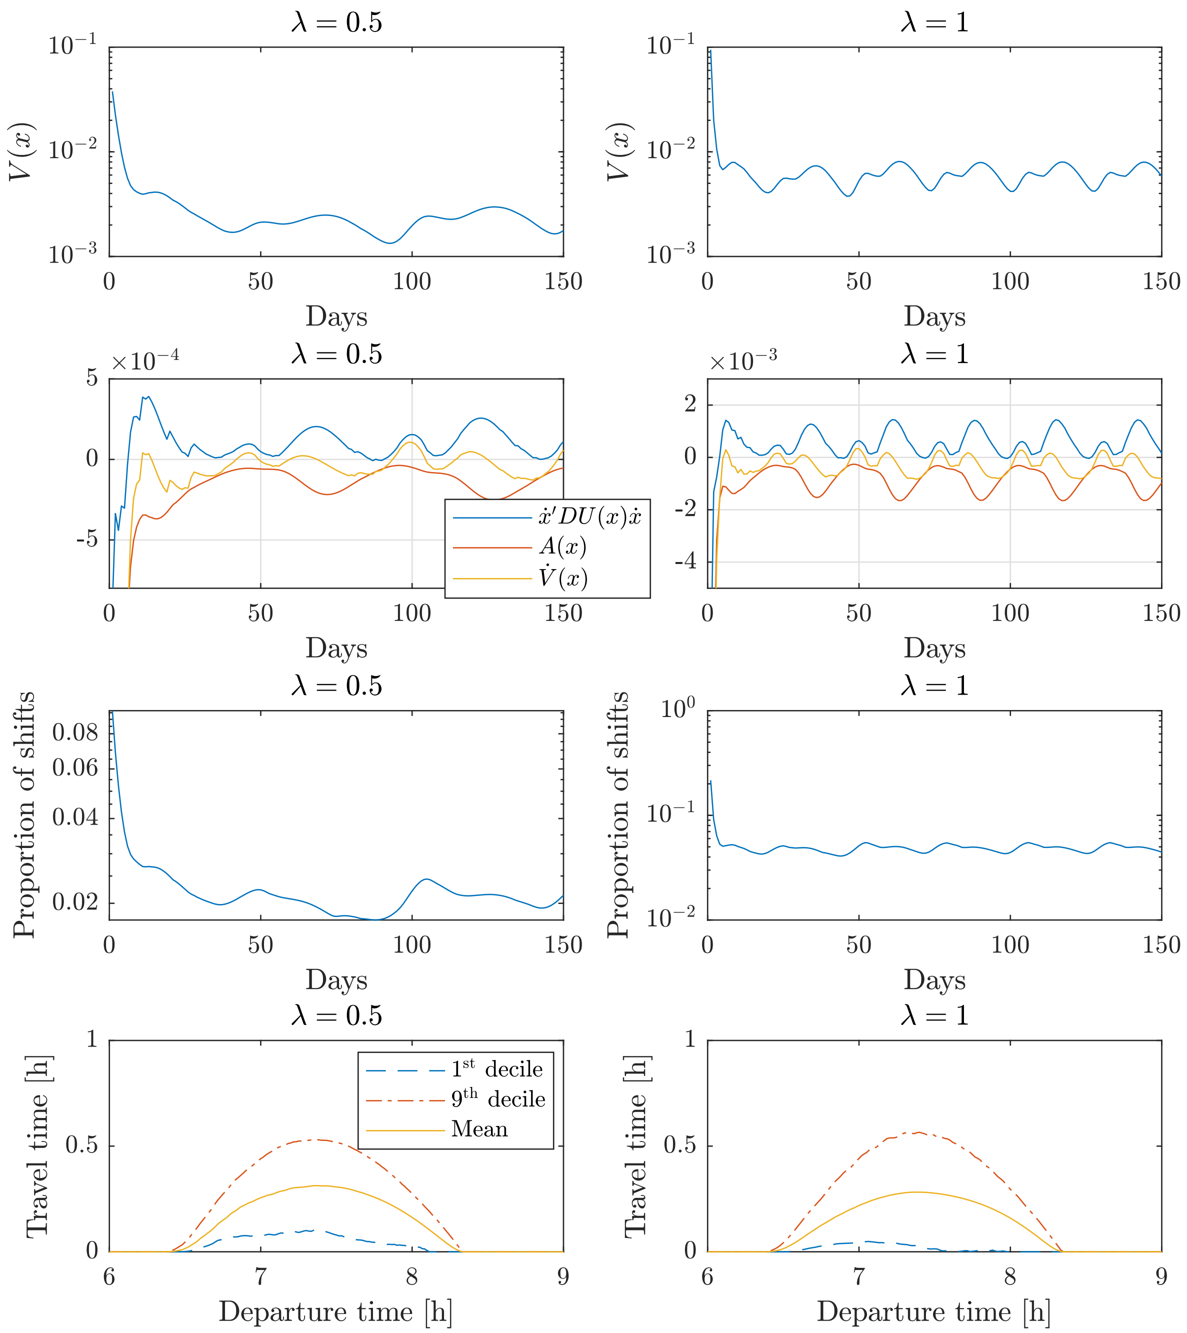

Figure 5: Two runs with homogeneous users ($t^* = 8$) and sensitivity $\lambda = 0.5$ and $\lambda = 1$.

3.3.2 Simulations

This section relies on simulation to complement the results previously derived analytically. The following simulations were carried out with the smooth preferences described in Section 3.2 and with the proportional swap adjustment mechanism. Additional details on the simulations are provided in appendix A.

Fig. 5 provides various graphs illustrating two runs of the dynamic process, with different sensitivity λ. The graphs in the first row represent the evolution of the function V, defined in Eq. (4) and providing an index of disequilibrium. The graphs in the second row represent the time derivative of $V(x)$ and its two components described in Eq. (3). The third row shows the proportion of users moved per day and the fourth describes the distribution of travel times over the last 100 days, as a function of departure time.

Note first that the term $\dot{x}DU(x)\dot{x}$ is almost always positive, suggesting that counter examples to monotonicity are the rule rather than exceptions. As expected, the term $A(x)$ is always negative. Its great magnitude during the first days brings the system close to equilibrium. Then, the magnitude of $A(x)$ decreases and $\dot{V}(x)$ becomes closer to zero. Interestingly however, $\dot{V}(x)$ remains mostly negative with both values of sensitivity, while $V(x)$ actually exhibits a quasi-periodic behavior. This apparent contradiction illustrates the limitations of using continuous-time tools for the study of a discrete-time systems. While the system is pushed by the adjustment mechanism in a direction that reduces the function $V(x)$ for small population shifts, larger shifts (obtained with a larger sensitivity λ) can make this distance increase. It is thus difficult to distinguish the effect of non-monotonicity from the effect of a too strong sensitivity. It would be interesting to determine whether there exists a threshold λ_0 such that reducing λ bellow λ_0 does not help reducing the oscillation size anymore. In practice, this threshold might be unrealistically small.

The last two rows provide some order of magnitude regarding the importance of this oscillatory behavior. On any given day, the proportion of the population that shifts departure time is about 2 % with $\lambda = 0.5$ and about 5 % with $\lambda = 1$.[3] This oscillatory behavior leads to a maximal difference of about 26 min between the 1^{st} and 9^{th} deciles of arrival time for $\lambda = 0.5$ and of about 32 min for $\lambda = 1$. Note that in the real world, daily variations in demand and supply produce additional variations in travel time uncertainty, producing a different dependance on departure time (Fosgerau, 2010).

[3]One may be tempted to conclude that these rather large proportions suggest an exceedingly large value of λ. Yet, we would like to highlight that these proportions largely depend on the state of the system itself and that with $\lambda = 0.5$, the system already takes about 20 days to approach its equilibrium. Given the rather high frequency of perturbations in the real world (e.g. seasonal demand and capacity variations), we believe that a system with a significantly lower sensitivity would not have enough time to approach its equilibrium before being perturbed again.

3.4 Heterogeneous users

3.4.1 Expected effects

We investigate in this section the influence of heterogeneity in schedule preferences on the size of oscillations by considering different t^*. Note that this can be interpreted either as heterogeneity in the marginal utility rate at work, or as heterogeneity in the constant travel time required between the bottleneck location and the working place. We consider heterogeneity in t^* because it allows breaking the symmetry of the system, thereby limiting the occurrence of unrealistic shifts from the congestion onset to the offset, and vice versa.

We are not aware of any tools that would allow us to study the stability of the morning commute with a continuum of heterogeneous users (with continuously distributed t^* here). Although we do not prove it, we expect the analytical monotonicity results derived in Section 3.3.1 to be transferable to the case with a finite number of homogeneous groups. We focus here on the effect of heterogeneity on the oscillation size, rather than on the existence of oscillations.

Intuitively, heterogeneity can be expected to reduce the oscillation size in various ways. At equilibrium, users are likely to strictly prefer their own departure time (or those of identical users) over those chosen by users with different preferences. This by itself is expected to reduce variability, as users should remain close to their equilibrium departures. Then, this reduction of the interval considered by each user for possible updates has other advantages. First, making shifts more local reduces the number of users that are perturbed by each decision update (recall that a user shifting from a departure time t to another t' only affects users that depart between t and t'). Second, all the adjustment mechanisms presented in Section 2.3 are based on the assumption that the probability to shift is proportional to the number of alternatives proposed (the distribution of utilities being the same). Thus, reducing the range of "interesting alternatives" reduces the shift frequency. Finally, making shifts more local also reduces the variations of marginal utility rate at work, which might contribute to curb the effects of non-monotonicity. The following section confronts these predictions with simulation results to quantify the effect of user heterogeneity. We do not attempt to distinguish the effects aforementioned as they are all intricately related.

3.4.2 Simulations

In the following simulations, the population is divided into N homogeneous groups, with only inter-group heterogeneity. Let x_i^p denote the proportion of the total population that belongs to group p and departs at time t_i, and $x_i = \sum_{p=1}^{N} x_i^p$.

We introduce two segregation metrics: group homogeneity is defined as the average over all users of the proportion of users belonging to their group among those that depart with them. It can be computed as $\sum_{i=1}^{n} \frac{\sum_{p=1}^{N}(x_i^p)^2}{x_i}$. We also define the w-diversity as $\sum_{i=1}^{n} x_i\sigma_i$, where σ_i is the standard deviation of the marginal rate at arrival among the users departing at t_i. Letting $a_i(x)$ denote the arrival time for these users, $\sigma_i = \sqrt{\sum_{p=1}^{N} \frac{x_i^p}{x_i}\left(w(a_i(x)) - \frac{\sum_{p=1}^{N} x_i^p w(a_i(x))}{x_i}\right)^2}$.

Before investigating the influence of heterogeneity itself, we ran multiple simulations with the same amount of heterogeneity, but with various numbers of groups (see the left column of Fig. 6). The values of t^* assigned to each group follow a discrete uniform distribution centered around 8 h with standard deviation $\sigma^* = 0.4$ h. It appears that the group homogeneity index is always significantly smaller than 1, indicating that users are for the most part rather far from being segregated. While the group homogeneity tends to increases as the disequilibrium index $V(x)$ decreases, it remains rather small while $V(x)$ reaches extremely small values (especially for large N). Note also that the w-diversity is essentially the same for $N = 10$ and $N = 20$. Altogether, this indicates that the system is rather insensitive to the exact number of groups, provided it is large enough.

The next simulations, illustrated in the second column of Fig. 6, were run with a given number of groups ($N = 10$) but various heterogeneity levels. Again, the values of t^* follow a discrete uniform distribution, but with different standard deviations σ^*. With small standard deviations, the differences in t^* are not sufficient to produce any segregation: the group homogeneity index is almost identical to the proportion that each group represents (0.1). The disequilibrium index $V(x)$ consequently follows an evolution that is extremely similar to the one observed with homogeneous users (see Fig. 5).

With larger σ^* however, the amplitude of oscillations decreases and oscillations fully disappear for $\sigma^* = 0.4$ h (for the time horizon considered). This phenomenon coincides with a large increase in group homogeneity. The rise in segregation between the cases $\sigma^* = 0.2$ h and 0.4 h is actually sufficient to keep the w-diversity constant at the 150$^{\text{th}}$ day, despite the additional population diversity.

Finally, Fig. 7 provides more intuitive orders of magnitude, that can be compared with those

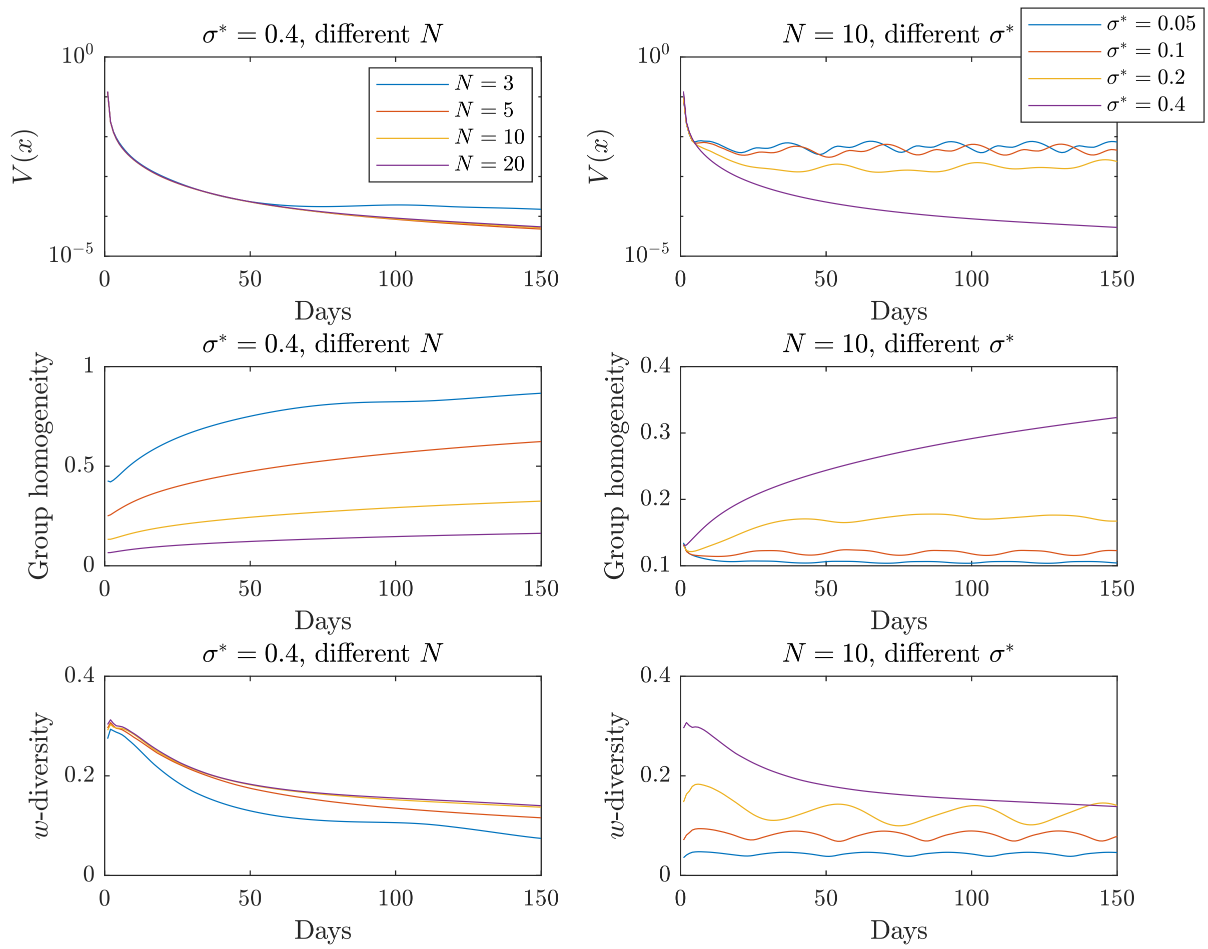

Figure 6: Simulations with various numbers of population groups (N) and various inter-group heterogeneity levels (σ^*) (sensitivity $\lambda = 1$).

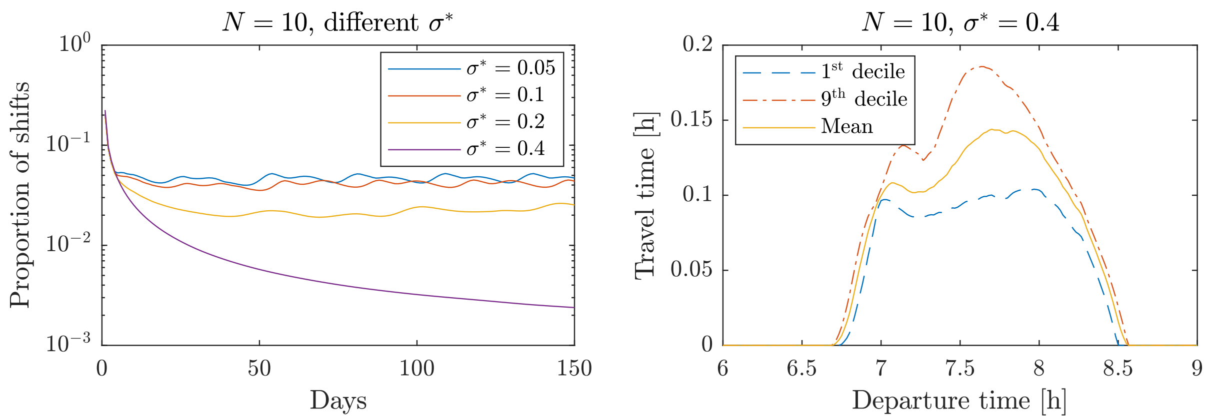

Figure 7: Impact of heterogeneity in t^* on variability.

provided in Fig. 5) for the homogeneous case. With $\sigma^* = 0.4$ h, the proportion of shifts per day decreases to less than 1% after about 28 days and the maximum difference between the 1st and the 9th decile of travel time over the last 100 days is approximately equal to 4 min.

4 Discussion

Our analysis has shown that the morning commute equilibrium should be expected to be unstable for a wide range of rational adjustment mechanisms and schedule preferences, due to the non-monotonicity of the utility function. This instability translates in a quasi-periodic oscillatory behavior. The magnitude of these oscillations and their proximity to the equilibrium depend strongly on the sensitivity of users to utility differences and on the amount of user heterogeneity.

Previous work (de Palma, 2000, Mc Breen *et al.,* 2006) identified similar dependence on user heterogeneity but could not fully explain the phenomena observed without the concept of monotonic utility functions. It is true that "congestion externality explains why several intuitive adjustment processes do not converge" (de Palma, 2000). Yet, continuous time stabity is not as much about the aggregate magnitude of externalities, as it is about their distribution across the other available alternatives. On-going work by the authors applying similar methods to another type of congestion mechanism (the Macroscopic Fundamental Diagram of an isotropic urban area) should provide further intuition regarding the complex dependency of stability on the distribution of externalities.

5 References

Arnott, R., A. de Palma and R. Lindsey (1988) Schedule delay and departure time decisions with heterogeneous commuters, *Transportation Research Record*, **1197**, 56–67.

Bressan, A., C. J. Liu, W. Shen and F. Yu (2012) Variational analysis of Nash equilibria for a model of traffic flow, *Quarterly of Applied Mathematics*, **70** (3) 495–515.

Daganzo, C. F. (1985) The uniqueness of a time-dependent equilibrium distribution of arrivals at a single bottleneck, *Transportation science*, **19** (1) 29–37.

de Palma, A. (2000) Solution and stability for a simple dynamic bottleneck model, in J. A. Filar, V. Gaitsgory and K. Mizukami (eds.) *Advances in Dynamic Games and Applications*, 405–425, Springer.

Fosgerau, M. (2010) On the relation between the mean and variance of delay in dynamic queues with random capacity and demand, *Journal of Economic Dynamics and Control*, **34** (4) 598–603, ISSN 0165-1889.

Fosgerau, M. and A. de Palma (2012) Congestion in a city with a central bottleneck, *Journal of Urban Economics*, **71** (3) 269–277.

Guo, R.-Y., H. Yang and H.-J. Huang (2018) Are we really solving the dynamic traffic equilibrium problem with a departure time choice?, *Transportation Science, Articles in Advance*, 1–18.

Hofbauer, J. and W. H. Sandholm (2009) Stable games and their dynamics, *Journal of Economic Theory*, **144** (4) 1665–1693.

Iryo, T. (2008) An analysis of instability in a departure time choice problem, *Journal of Advanced Transportation*, **42** (3) 333–358.

Lindsey, R. (2004) Existence, uniqueness, and trip cost function properties of user equilibrium in the bottleneck model with multiple user classes, *Transportation Science*, **38** (3) 293–314.

Maynard Smith, J. and G. R. Price (1973) The logic of animal conflict, *Nature*, **246**, 15–18.

Mc Breen, J., P. Jensen and F. Marchal (2006) An agent based simulation model of traffic congestion, paper presented at the *ATT06: Proceedings of the 4th workshop on Agents in Traffic and Transportation*, 43–49.

Mounce, R. (2006) Convergence in a continuous dynamic queueing model for traffic networks, *Transportation Research Part B: Methodological*, **40** (9) 779–791.

Newell, G. F. (1987) The morning commute for nonidentical travelers, *Transportation Science*, **21** (2) 74–88.

Sandholm, W. H. (2010a) Local stability under evolutionary game dynamics, *Theoretical Economics*, **5**, 27–50.

Sandholm, W. H. (2010b) *Population games and evolutionary dynamics*, MIT press.

Sandholm, W. H. (2015) Chapter 13 - population games and deterministic evolutionary dynamics, in H. P. Young and S. Zamir (eds.) *Handbook of Game Theory with Economic Applications*, vol. 4, 703–778, Elsevier.

Small, K. A. (2015) The bottleneck model: An assessment and interpretation, *Economics of Transportation*, **4**, 110–117. Special Issue on Collective Contributions in the Honor of Richard Arnott.

Smith, M. and M. Ghali (1990) The dynamics of traffic assignment and traffic control: A theoretical study, *Transportation Research Part B: Methodological*, **24** (6) 409–422.

Smith, M. J. (1984a) The existence of a time-dependent equilibrium distribution of arrivals at a single bottleneck, *Transportation Science*, **18** (4) 385–394.

Smith, M. J. (1984b) The stability of a dynamic model of traffic assignment: An application of a method of lyapunov, *Transportation Science*, **18** (3) 245–252.

Taylor, P. D. and L. B. Jonker (1978) Evolutionary stable strategies and game dynamics, *Mathematical biosciences*, **40** (1-2) 145–156.

Tseng, Y.-Y. and E. T. Verhoef (2008) Value of time by time of day: A stated-preference study, *Transportation Research Part B: Methodological*, **42** (7?8) 607 – 618, ISSN 0191-2615.

Vickrey, W. (1973) Pricing, metering, and efficiently using urban transportation facilities, *Highway Research Record*, **476**, 36–48.

A Implementation of the simulator

The simulations were carried out for a case where it takes 2h to serve all users, with possible departures regulary spaced every 1 min from 6 AM to 9 AM.

To better represent the real world, the simulator was designed such that users choosing to depart at time t_i were actually departing uniformly over the interval $[t_i - \frac{\delta t}{2}, t_i + \frac{\delta t}{2}]$, where $\delta t = 1min$ is the time between two consecutive departure time alternatives. The resulting departure rate is piece-wise constant, allowing for simple analytical derivations of the queue dynamics. The experienced utility was then computed based on the distribution of arrival times for users departing $[t_i - \frac{\delta t}{2}, t_i + \frac{\delta t}{2}]$.