

Providing slowdown information to improve selfish routing [★]

Philip N. Brown¹[0000–0003–3953–0503]

University of Colorado Colorado Springs, CO, USA
`philip.brown@uccs.edu`

Abstract. Recent research in the social sciences has identified situations in which small changes in the way that information is provided to consumers can have large aggregate effects on behavior. This has been promoted in popular media in areas of public health and wellness, but its application to other areas has not been broadly studied. This paper presents a simple model which expresses the effect of providing commuters with carefully-curated information regarding aggregate traffic “slowdowns” on the various roads in a transportation network. Much of the work on providing information to commuters focuses specifically on travel-time information. However, the model in the present paper allows a system planner to provide slowdown information as well; that is, commuters are additionally told how much slower each route is as compared to its uncongested state. We show that providing this additional information can improve equilibrium routing efficiency when compared to the case when commuters are only given information about travel time, but that these improvements in congestion are not universal. That is, transportation networks exist on which any provision of slowdown information can harm equilibrium congestion. In addition, this paper illuminates a deep connection between the effects of commuter slowdown-sensitivity and the study of marginal-cost pricing and altruism in congestion games.

Keywords: Congestion Game · Traffic Information System · Transportation Networks.

1 Introduction

Today, in the age of the internet of things, big data, and pervasive computing, engineered systems are becoming increasingly interconnected with the human populations that they serve. System-level performance is directly affected by the choices of human users, customers, and adversaries, and it has long been recognized that self-interested behavior by users may lead to gross system inefficiencies [21, 22]. Thus, both the importance and the feasibility of influencing human behavior in intelligent ways are increasing simultaneously.

A popular modeling framework to study methods of influencing selfish behavior in engineered systems is the *nonatomic congestion game*, which is often

[★] This work is supported in part by NSF Grant ECCS-2013779.

used to model urban transportation networks [2]. This model was one of the first to admit straightforward characterizations of the *price of anarchy*, a popular measure of the social cost of selfish behavior [22]. It is also a natural setting to study various modified models of human decision-making; examples include altruism [10], pessimism [19], risk-aversion [17], and various generalizations [15]. More relevant to this paper, it has proved fertile ground for the study of various methods of influencing user behavior: taxation [4, 7, 12, 13], autonomously-controlled traffic [3, 16], and traffic information systems [1, 11, 18] have all been investigated in this context. The majority of this literature depends on the usual game-theoretic assumption that users are expected utility maximizers.

In parallel, an increasingly popular conceptual framework for studying the practice of influencing human behavior is that of *behavioral economics*. This framework grew out of empirical studies which showed that humans tend to deviate from the behaviors prescribed by expected utility maximization, and that these deviations manifest themselves in predictable ways [25]. If humans are predictably irrational, it seems reasonable that engineers might attempt to exploit this predictability to influence behavior. This idea has led to the concept of the “nudge” — the popular name given to the idea that making slight modifications to a person’s environment can have a profound effect on their behavior [24].

This paper draws inspiration from the concept of the nudge, and asks if drivers in transportation networks could perhaps be influenced by providing simple pieces of information that are carefully designed to exploit behavioral biases in their decision-making process. As a simple preliminary model for studying such questions, this paper begins with the typical assumption that drivers prefer low travel times over high travel times. Some drivers, however, additionally dislike the idea of choosing a route that exhibits higher travel time than its uncongested free-flow travel time. That is, some drivers are slowdown-averse. A system planner, knowing about the existence (but perhaps not the magnitude) of this slowdown-aversion, may wish to exploit it to improve aggregate congestion, and provides a signal informing drivers of the presence and magnitude of slowdowns on various routes. Note that unlike much other work on traffic information systems, we do not require ordinary drivers to perform any type of Bayesian inference; rather, our model captures a phenomenon in which each driver acts in response to her “gut feel” about how much she dislikes traffic slowdowns.

The main question of this paper is this: when can a planner be certain that providing slowdown information to drivers will improve congestion? In Theorem 1, we show that networks exist for which any slowdown information provided to drivers increases equilibrium delays, and thus on these networks, the planner’s only reasonable course of action is to provide no slowdown signal at all. Nonetheless, Theorem 1 also shows that on parallel networks, slowdown information can always reduce equilibrium delays. Perhaps unintuitively, these results are tightly connected with the literature on marginal-cost pricing and altruistic behavior in transportation networks [5, 10]. Indeed, this condition can be stated as a sufficient condition for improvements due to slowdown signaling: If altru-

ism increases equilibrium delays on some network, then slowdown signaling also increases equilibrium delays on that network.

Following this, Theorem 2 essentially asks the question “is it optimal for planners to provide accurate slowdown information to drivers?” Alternatively, what is the optimal slowdown signal with respect to minimizing equilibrium congestion? Here, we consider an example setting in which the planner has significant levels of information about the population’s slowdown preferences, as well as some minimal information about aggregate demand. In this setting, Theorem 2 shows that planners with access to this additional information optimally *overstate* the severity of network slowdowns. That is, a planner’s optimal course of action may well be to deliberately provide misleading information to drivers.

Lastly, the paper closes with a note on the implications that this work has for the broader study of heterogeneous nonatomic congestion games. Here, it is shown that many of the well-studied forms of heterogeneous congestion games appearing in the literature belong to a particularly well-behaved class of games known as *weighted potential games*. That is, the Nash equilibria of these games can be characterized as the maximizers of a single global convex potential function. From an analytical point of view, the ramifications of this are broad: the existing work on these types of games depends on *ad hoc* characterizations of equilibrium behavior, and the existence of a potential function may unify these approaches.

2 Model and Performance Metrics

2.1 Routing Game

Consider a network routing problem for a network (V, E) comprised of vertex set V and edge set E . We call a source/destination vertex pair $(\sigma^c, t^c) \in (V \times V)$ a *commodity*, and the set of all such commodities $\mathcal{C}$. For each commodity $c \in \mathcal{C}$, there is a mass of traffic $r^c > 0$ that needs to be routed from σ^c to t^c . We write $\mathcal{P}^c \subset 2^E$ to denote the set of *paths* available to traffic in commodity c , where each path $p \in \mathcal{P}^c$ consists of a set of edges connecting σ^c to t^c . Let $\mathcal{P} = \cup \{\mathcal{P}^c\}$. A network is called *symmetric* if there is exactly one commodity: $\mathcal{C} = \{c\}$, i.e., all traffic routes from a common source σ to a common destination t using a common path set $\mathcal{P}$. A network is called a *parallel* network if all commodities share a single source-destination pair and all paths are disjoint; i.e., for all paths $p, p' \in \mathcal{P}$, $p \cap p' = \emptyset$. Note that a parallel network need not be symmetric; although all traffic must share a common source and destination, the path sets $\mathcal{P}^c$ available to traffic from different commodities may differ.

We write $f_p^c \geq 0$ to denote the mass of traffic from commodity c using path p , and $f_p := \sum_{c \in \mathcal{C}} f_p^c$. A *feasible flow* $f \in \mathbb{R}^{|\mathcal{P}|}$ is an assignment of traffic to various paths such that for each c , $\sum_{p \in \mathcal{P}^c} f_p^c = r^c$ and $\sum_{c \in \mathcal{C}} r^c = r$.

Given a flow f , the flow on edge e is given by $f_e = \sum_{p: e \in p} f_p$. To characterize transit delay as a function of traffic flow, each edge $e \in E$ is associated with a specific latency function $\ell_e : [0, r] \rightarrow [0, \infty)$; $\ell_e(f_e)$ denotes the delay

experienced by users of edge e when the edge flow is f_e . We adopt the standard assumptions that each latency function is nondecreasing, convex, and continuously differentiable. We measure the cost of a flow f by the *total latency*, given by

$$\mathcal{L}(f) = \sum_{e \in E} f_e \cdot \ell_e(f_e) = \sum_{p \in \mathcal{P}} f_p \cdot \ell_p(f), \quad (1)$$

where $\ell_p(f) = \sum_{e \in p} \ell_e(f_e)$ denotes the latency on path p . We denote the flow that minimizes the total latency by

$$f^* \in \underset{f \text{ is feasible}}{\operatorname{argmin}} \mathcal{L}(f). \quad (2)$$

A *routing problem* is given by $G = (V, E, \mathcal{C}, \{\ell_e\})$. Classes of routing problems are denoted by $\mathcal{G}$. For $d \geq 1$, we write $\mathcal{G}^d$ to denote the class of all routing problems with latency functions of the form $\ell_e(f_e) = a_e(f_e)^d + b_e$, where $a_e, b_e \geq 0$ are edge-specific constants.

To study the effect of slowdown information on self-interested behavior, this paper models the above routing problem as a heterogeneous non-atomic congestion game. The slowdown-sensitivities of the users in commodity c are modeled by a monotone, nondecreasing function $\beta^c : [0, r^c] \rightarrow [0, 1]$, where each user $x \in [0, r^c]$ has a slowdown sensitivity $\beta^{c,x} \in [0, 1]$. Here, if user x has $\beta^{c,x} = 0$, this indicates that the user is purely delay-averse and is unresponsive to slowdowns. On the other hand, if user x has $\beta^{c,x} = 1$, this indicates that the user cares nothing for *actual* delay, and selects routes solely on the basis of how over-congested the route is compared to nominal. The analysis in this paper assumes that each sensitivity distribution function β^c is unknown to the system operator, and we write $\mathcal{B}$ to denote the set of all feasible sensitivity distribution functions

The system operator provides all users with information regarding the slowdown experienced on the various paths; this is modeled by a number $\gamma \in [0, 1]$. When $\gamma = 1$, users are provided complete and true information regarding each road's slowdown; when $\gamma = 0$, users are provided no information regarding slowdowns. Altogether, given a flow f , the subjective cost that user $x \in [0, r^c]$ experiences for using path $p \in \mathcal{P}^c$ is of the form

$$\begin{aligned} J^{c,x}(p; f) &= \sum_{e \in p} [(1 - \gamma\beta^{c,x}) \ell_e(f_e) + \gamma\beta^{c,x} (\ell_e(f_e) - \ell_e(0))] \\ &= \sum_{e \in p} [\ell_e(f_e) - \gamma\beta^{c,x} \ell_e(0)]. \end{aligned} \quad (3)$$

Thus, each user $x \in [0, r^c]$ can be viewed as interpreting the cost of a road as the difference between its latency and a moderately-scaled version of its free-flow delay. We assume that each user selects the lowest-cost path from the available source-destination paths. We call a flow f a *Nash flow* if all users are individually using minimum-cost paths given the choices of other users. That is, for all commodities $c \in \mathcal{C}$, every user $x \in [0, r^c]$ using path p in f experiences a cost satisfying

$$J^{c,x}(p; f) = \min_{\tilde{p} \in \mathcal{P}^c} J^{c,x}(\tilde{p}; f). \quad (4)$$

The above game is an *exact potential game* with a convex potential function (see Lemma 2), which implies that its set of Nash flows is nonempty and convex.

2.2 Avoiding perverse signaling

As an initial step towards characterizing a system operator’s optimal signaling policy, this paper compares the equilibrium total latency resulting from slowdown signaling with the un-influenced equilibrium total latency. One simple measure which captures the possible *harm* of slowdown signaling is the *perverseity index* as introduced in [9]:

$$\text{PI}(\mathcal{G}, \gamma) := \sup_{G \in \mathcal{G}} \sup_{\beta \in \mathcal{B}} \frac{\mathcal{L}^{\text{nf}}(G, \beta, \gamma)}{\mathcal{L}^{\text{nf}}(G, \beta, 0)}. \quad (5)$$

Here, if a class of routing problems $\mathcal{G}$ and signal parameter γ have a large perverseity index, this indicates that there exist routing problems in $\mathcal{G}$ for which it would be better for the system planner to avoid signaling altogether. If $\text{PI}(\mathcal{G}, \gamma) > 1$ for some γ , we say that γ is a *perverse* signaling policy. If $\text{PI}(\mathcal{G}, \gamma) = 1$, we say that γ is *non-perverse*.

3 Related Work

Since the seminal work on price of anarchy for nonatomic routing games [22], these games have provided fertile ground to investigate a wide range of questions. A common theme in these involves augmenting users’ costs in some way to investigate the effect of some external influence or internal bias. For example, the altruism models of [5, 10] assume that each member $x \in [0, 1]$ of the population has an “altruism” parameter $\alpha^x \in [0, 1]$ which enters their edge cost function in the following way:

$$J_e^x(f_e) = \ell_e(f_e) + \alpha^x f_e \ell'_e(f_e). \quad (6)$$

Here, $\alpha^x = 1$ represents a fully altruistic user whose singular goal is to reduce aggregate congestion, and $\alpha^x = 0$ represents a fully selfish user whose singular goal is to minimize personal travel time. This model of altruism has been studied extensively, and much is known about the effects of this type of altruistic bias — notably, in many situations, increased levels of altruism lead to reductions in equilibrium delays. Many other similar cost function biases have been studied which take a similar form of $\ell_e(f_e) + b_e(f_e)$ for some specified $b_e(f_e)$, including pessimism [19], risk-aversion [17], and various generalizations [15, 20].

The above work is concerned with *characterizing* the effect of user preferences on equilibrium behavior. In parallel, many researchers have investigated methods of *influencing* user preferences to improve equilibrium behavior. These works on influence often involve similar modifications of user costs. For example, pricing [14] is a common means of influencing user choices, and a popular pricing model allows the system planner to assign pricing functions to each edge

of $\tau_e(f_e)$, and assumes (similar to the altruism model above) that users have heterogeneous price-sensitivities s^x :

$$J_e^x(f_e) = \ell_e(f_e) + s^x \tau_e(f_e). \quad (7)$$

This is clearly tightly connected with the altruism model, particularly if the $\{\tau_e\}$ are chosen to be marginal-cost prices of the form $\tau_e^{\text{mc}}(f_e) = f_e \ell'_e(f_e)$. This has led to various synergies between the study of altruism and marginal-cost pricing [5], as results in one area neatly imply conclusions in the other.

4 Our Contributions on Slowdown Signaling

4.1 Providing slowdown information can harm congestion on some networks

Our first result shows that no universal signaling policy exists: networks exist for which all nonzero slowdown signals can increase aggregate congestion costs. However, some classes of networks are immune to these pathologies; in particular, slowdown signaling has the potential to improve congestion on all parallel networks. In the following, for simplicity of exposition, each edge's latency function is assumed to be a polynomial of the form $\ell_e(f_e) = a_e(f_e)^d + b_e$, for edge-specific nonnegative coefficients a_e and b_e .

Theorem 1. *Let $\mathcal{G}$ be the class of all routing problems. For all $\gamma \in (0, 1]$, routing problems exist for which γ is a perverse signal:*

$$\text{PI}(\mathcal{G}, \gamma) > 1. \quad (8)$$

However, let $\mathcal{G}_p^d$ be the class of all parallel routing problems with latency functions of the form $\ell_e(f_e) = a_e(f_e)^d + b_e$. Then non-perverse signaling is possible. In particular,

$$\gamma \in [0, d/(d+1)] \quad \text{if and only if} \quad \text{PI}(\mathcal{G}_p^d, \gamma) = 1. \quad (9)$$

Note that since $\mathcal{G}_p^d \subset \mathcal{G}$, the results in (8) and (9) indicate that there exist networks for which all signaling policies are perverse, but that these pathological networks are *never* parallel networks. Thus, a system planner who knows they are working with parallel networks can safely employ any signal satisfying $\gamma \in [0, d/(d+1)]$ without fear of causing harm relative to the uninfluenced equilibria.

The proof of this theorem relies on Lemma 1, which explicitly relates the equilibrium flows under slowdown signaling to a similar formulation which is reminiscent of user costs experienced under marginal-cost tolls. This allows us to leverage existing results on marginal-cost tolls to obtain the proof of the theorem.

Lemma 1. *Let $\beta \in \mathcal{B}$, $G \in \mathcal{G}^d$, $\gamma \in [0, 1)$. For each $x \in [0, 1]$, let*

$$\alpha^x := \frac{\gamma \beta^x}{d(1 - \gamma \beta^x)}. \quad (10)$$

Then the set of Nash flows associated with the game (G, β, γ) is equal to the set of Nash flows for a game with the same network and a user-specific edge cost function of

$$J_e^x(f_e) = (1 + d\alpha^x)a_e(f_e)^d + b_e. \quad (11)$$

The proof of Lemma 1 is provided in the Appendix. The cost functions (11) in Lemma 1 are well-studied in the literature on heterogeneous nonatomic routing games. These cost functions are identical to those induced by marginal-cost tolls when users are heterogeneous in price sensitivity; they also appear in the α -altruism model of [5, 10]. Thus, results from those two streams of work may be leveraged to draw conclusions about the effects of slowdown signaling.

Proof of Theorem 1 Due to Lemma 1, the proof of Theorem 1 follows in a straightforward manner from Theorem (2) of [9] and Theorem 7.1 of [10]. $\square$

4.2 Optimal signaling need not be truthful

Our next result considers the goal of computing the optimal signaling policy which minimizes worst-case equilibrium traffic congestion. Here, we consider a situation in which the planner possesses additional information and in which case aggressive signaling mechanisms may be warranted. One example of this is reminiscent of the setting studied in [8]: the planner knows *a priori* that the traffic rate on the network is high enough that all edges would be used in an un-influenced Nash flow. Here, suppose that the planner knows that for every user $x \in [0, 1]$, the slowdown-sensitivity satisfies $\beta^x \in [\beta_L, \beta_U]$, where $\beta_L > 0$ and $\beta_U < 1$. In this setting, for linear-affine-cost parallel networks, the planner's optimal signal may actually *over-state* the network's true slowdowns:

Theorem 2. *Let $\bar{\mathcal{G}}_p^1$ denote the class of linear-latency parallel networks in which every edge has positive traffic in an un-influenced Nash flow. For every $G \in \bar{\mathcal{G}}_p^1$, it holds that*

$$\gamma^* := \frac{1}{\beta_L + \beta_U} = \arg \min_{\gamma \geq 0} \max_{\beta} \mathcal{L}^{\text{nf}}(G, \beta, \gamma). \quad (12)$$

Under the influence of this signal, the worst-case equilibrium total latency satisfies

$$\max_{G \in \bar{\mathcal{G}}_p^1} \max_{\beta} \frac{\mathcal{L}^{\text{nf}}(G, \beta, \gamma^*)}{\mathcal{L}^*(G)} = \frac{4}{3} \left(1 - \frac{\beta_L/\beta_U}{(1 + \beta_L/\beta_U)^2} \right). \quad (13)$$

The proof of Theorem 2 is provided in the Appendix.

4.3 Broader Implications for Heterogeneous Congestion Games

An interesting byproduct of the analysis required for this work is that we have discovered a weighted potential game formulation which applies to a large class of *heterogeneous* nonatomic congestion games; to the best of our knowledge, this formulation is novel. This formulation opens the door to computing Nash flows efficiently for a large class of games by performing gradient descent on the potential function of an associated exact potential game. To understand the new formulation, first consider the following lemma which demonstrates that the slowdown-sensitivity games in this paper are exact potential games:

Lemma 2. *For any $\beta \in \mathcal{B}$, $G \in \mathcal{G}$, and $\gamma \geq 0$, the game (G, β, γ) specified by cost functions (3) is an exact potential game with a convex potential function. Thus:*

1. (G, β, γ) has at least one Nash flow,
2. the set of Nash flows is convex, and
3. for each user $x \in [0, 1]$, given two Nash flows f and f' (in which user x chooses paths p and p' , respectively), it holds that $J^x(p; f) = J^x(p'; f')$.

The proof of Lemma 2 is provided in the Appendix. We can now state the main result, that many heterogeneous nonatomic congestion games with a particular form of polynomial cost function are weighted potential games and that their Nash flows exactly coincide with those of our slowdown-sensitivity games.

Theorem 3. *In some routing problem $G \in \mathcal{G}^d$, let every user $x \in [0, 1]$ have altruism parameter $\alpha^x \in [0, 1]$ and experience a cost on edge e of $J_e^x(f) = (1 + d\alpha^x)a_e(f_e)^d + b_e$. Then routing problem G coupled with the user population described by α^x is a weighted potential game whose set of Nash flows is equal to that of slowdown-sensitivity game $(G, \beta, 1)$, where for each $x \in [0, 1]$ it holds that*

$$\beta^x = \frac{\alpha^x d}{\alpha^x d + 1}. \quad (14)$$

The proof of Theorem 3 is provided in the Appendix.

5 Conclusion

This paper explores the connections between a new class of slowdown-sensitive congestion games and established results for heterogeneous nonatomic routing games under the influence of altruism and marginal-cost pricing. In particular, we demonstrate a tight connection between the two, and illustrate how to connect results from one to the other. Finally, we exploit this connection to demonstrate a novel result for certain heterogeneous nonatomic congestion games, namely that they are weighted potential games and thus their equilibria may be computed efficiently.

References

1. Acemoglu, D., Makhdoumi, A., Malekian, A., Ozdaglar, A.: Informational Braess' Paradox: The Effect of Information on Traffic Congestion. *Operations Research* **66**(4) (2018)
2. Beckmann, M.J., McGuire, C.B., Winsten, C.B.: *Studies in the Economics of Transportation* (1955). <https://doi.org/10.1057/jors.1980.83>
3. Biyik, E., Lazar, D., Pedarsani, R., Sadigh, D.: Incentivizing Efficient Equilibria in Traffic Networks with Mixed Autonomy. *IEEE Transactions on Control of Network Systems* pp. 1–1 (2021). <https://doi.org/10.1109/TCNS.2021.3084045>, <https://ieeexplore.ieee.org/document/9442301/>
4. Bonifaci, V., Salek, M., Schäfer, G.: Efficiency of Restricted Tolls in Non-atomic Network Routing Games. In: *Algorithmic Game Theory*, vol. 6982 LNCS, pp. 302–313 (2011), http://link.springer.com/10.1007/978-3-642-24829-0_27
5. Brown, P.N.: When Altruism is Worse than Anarchy in Nonatomic Congestion Games. In: *2021 American Control Conference (ACC)*. pp. 4503–4508. IEEE (may 2021). <https://doi.org/10.23919/ACC50511.2021.9483340>, <http://arxiv.org/abs/2007.05591> <https://ieeexplore.ieee.org/document/9483340/>
6. Brown, P.N., Marden, J.R.: A Study on Price-Discrimination for Robust Social Coordination. In: *American Control Conf.* pp. 1699–1704. Boston, MA, USA (2016)
7. Brown, P.N., Marden, J.R.: Studies on Robust Social Influence Mechanisms: Incentives for Efficient Network Routing in Uncertain Settings. *IEEE Control Systems Magazine* **37**(1), 98–115 (2017). <https://doi.org/10.1109/MCS.2016.2621482>
8. Brown, P.N., Marden, J.R.: The Robustness of Marginal-Cost Taxes in Affine Congestion Games. *Transactions on Automatic Control* **62**(8), 3999–4004 (2017)
9. Brown, P.N., Marden, J.R.: Can Taxes Improve Congestion on All Networks? *IEEE Transactions on Control of Network Systems* **7**(4), 1643–1653 (2020). <https://doi.org/10.1109/TCNS.2020.2992679>
10. Chen, P.a., De Keijzer, B., Kempe, D., Shaefer, G.: Altruism and Its Impact on the Price of Anarchy. *ACM Trans. Economics and Computation* **2**(4) (2014)
11. Cheng, Y., Langbort, C.: A model of informational nudging in transportation networks. In: *2016 IEEE 55th Conference on Decision and Control, CDC 2016*. pp. 7598–7604 (2016). <https://doi.org/10.1109/CDC.2016.7799443>
12. Cole, R., Dodis, Y., Roughgarden, T.: Pricing network edges for heterogeneous selfish users. In: *Proc. of the 35th ACM symposium on Theory of computing*. pp. 521–530. New York, New York, USA (2003). <https://doi.org/10.1145/780615.780618>, <http://portal.acm.org/citation.cfm?doid=780542.780618>
13. Colini-Baldeschi, R., Klimm, M., Scarsini, M.: Demand-Independent Tolls pp. 1–18 (2017), <http://arxiv.org/abs/1708.02737>
14. Ferguson, B.L., Brown, P.N., Marden, J.R.: The Effectiveness of Subsidies and Tolls in Congestion Games. *IEEE Transactions on Automatic Control* pp. 1–1 (feb 2021). <https://doi.org/10.1109/TAC.2021.3088412>, <http://arxiv.org/abs/2102.09655> <https://ieeexplore.ieee.org/document/9451652/>
15. Kleer, P., Schäfer, G.: Path deviations outperform approximate stability in heterogeneous congestion games. In: *International Symposium on Algorithmic Game Theory*. pp. 212–224 (2017), <http://arxiv.org/abs/1707.01278>
16. Li, R., Brown, P.N., Horowitz, R.: Employing Altruistic Vehicles at On-ramps to Improve the Social Traffic Conditions. In: *2021 American Control Conference (ACC)*. pp. 4547–4552. IEEE (may 2021). <https://doi.org/10.23919/ACC50511.2021.9482993>, <https://ieeexplore.ieee.org/document/9482993/>

17. Lianas, T., Nikolova, E., Stier-Moses, N.E.: Asymptotically tight bounds for inefficiency in risk-averse selfish routing. In: IJCAI International Joint Conference on Artificial Intelligence. vol. 2016-Janua, pp. 338–344 (2016)
18. Massicot, O., Langbort, C.: Public Signals and Persuasion for Road Network Congestion Games under Vagaries. IFAC-PapersOnLine **51**(34), 124–130 (2019). <https://doi.org/10.1016/j.ifacol.2019.01.054>, <https://doi.org/10.1016/j.ifacol.2019.01.054>
19. Meir, R., Parkes, D.C.: Playing the Wrong Game: Smoothness Bounds for Congestion Games with Behavioral Biases. Performance Evaluation Review **43**(3), 67–70 (2015). <https://doi.org/10.1145/2847220.2847242>, <http://arxiv.org/abs/1411.1751>
20. Meir, R., Parkes, D.C.: Playing the Wrong Game: Bounding Negative Externalities in Diverse Populations of Agents. In: 17th International Conference on Autonomous Agents and MultiAgent Systems. pp. 86–94 (2018), <https://arxiv.org/pdf/1411.1751.pdf>
21. Pigou, A.C.: The Economics of Welfare. Macmillan, New York (1920)
22. Roughgarden, T.: Selfish Routing and the Price of Anarchy. MIT Press (2005)
23. Sandholm, W.H.: Population Games and Evolutionary Dynamics. MIT Press (2009)
24. Thaler, R.H., Sunstein, C.R.: Nudge (2008)
25. Tversky, A., Khaneman, D.: Judgment under Uncertainty: Heuristics and Biases. Science (1974)

Appendix: Proofs

Proof of Lemma 1. By (10), emulated cost functions (11) can be written

$$J_e^x(f_e) = \frac{a_e(f_e)^d}{1 - \gamma\beta^x} + b_e. \quad (15)$$

Since users’ ordinal preferences (and thus Nash flows) are invariant to multiplication by user-specific constants, the cost functions in (15) are equivalent to ones given by

$$\tilde{J}_e^x(f_e) = a_e(f_e)^d + (1 - \gamma\beta^x)b_e. \quad (16)$$

Evidently, the cost functions (16) are equal to those of the nominal game (G, β, γ) (see (3)). Since, for each user x , these cost functions encode the same ordinal preferences as those given by (11), the two sets of cost functions induce identical sets of Nash flows. $\square$

Proof of Theorem 2. The optimal signal factor (12) can be deduced from the results in [8, Theorem 1] regarding the optimal marginal-cost-toll scale factor in an identical setting. Our Lemma 1 provides that once a value of γ is fixed, each slowdown-sensitivity β^x distribution can be used to compute a taxation-sensitivity α^x distribution which induces an identical set of Nash flows. In [8, Theorem 1], it is shown that worst-case congestion is minimized for scaled marginal-cost taxes when the taxation-sensitivity distribution satisfies

$$\max_x \alpha^x = \frac{1}{\min_x \alpha^x}. \quad (17)$$

Thus, due to the equivalence between slowdown-sensitivity and taxation sensitivity and applying the transformation (10) from Lemma 1, the optimal signal factor γ^* satisfies

$$\frac{\gamma^* \beta_L}{1 - \gamma^* \beta_L} = \frac{\gamma^* \beta_U}{1 - \gamma^* \beta_U}, \quad (18)$$

which implies (12). The upper bound (13) follows immediately from [6, Lemma 3.1]. $\square$

Proof of Lemma 2.

Note that the individual slowdown sensitive cost function (3) is a sum of two terms: $\ell_e(f_e)$ and $\gamma \beta^x \ell_e(0)$. A game with costs defined by only the first term is simply a homogeneous nonatomic routing game, which is well-known to be an exact potential game with a convex potential function [22, 23]. A game with costs defined by only the second term is a trivial game in which each user's cost function depends only on her own action; any such game is also known to be an exact potential game with a linear potential function [23]. In [23] it is shown that a game whose cost functions are the sum of exact potential games' cost functions is itself an exact potential game whose potential function is the sum of the potential functions of its component games. Therefore, (G, β, γ) is an exact potential game with a convex potential function — and this convexity implies points (1), (2), and (3). $\square$

Proof of Theorem 3.

The proof relies on a serial application of Lemmas 1 and 2. In particular, Lemma 1 provides a bijection between the heterogeneous altruistic congestion games assumed by Theorem 3 and the slowdown-sensitivity games considered in this paper. Equation (14) is simply the inverse mapping of (10); thus, the cost functions of the two games express identical ordinal preferences and are related by a set of user-specific constant multipliers. Thus, the games are (for the purposes of equilibrium computation) equivalent: equilibria computed for one are automatically equilibria for the other. Accordingly, in light of Lemma 2, one may use potential game equilibrium-finding techniques (e.g., gradient descent) to compute the equilibria of the slowdown-sensitivity game and thereby find the equilibria of the altruistic congestion game. $\square$