

When Does Diversity of Agent Preferences Improve Outcomes in Selfish Routing?

Richard Cole¹, Thanasis Lianas² and Evdokia Nikolova³

¹ New York University,

² National Technical University of Athens,

³ University of Texas at Austin,

cole@cs.nyu.edu, lianas@corelab.ntua.gr, nikolova@austin.utexas.edu

Abstract

We seek to understand when heterogeneity in agent preferences yields improved outcomes in terms of overall cost. That this might be hoped for is based on the common belief that diversity is advantageous in many multi-agent settings. We investigate this in the context of routing. Our main result is a sharp characterization of the network settings in which diversity always helps, versus those in which it is sometimes harmful.

Specifically, we consider routing games, where diversity arises in the way that agents trade-off two criteria (such as time and money, or, in the case of stochastic delays, expectation and variance of delay). Our main contributions are: 1) A participant-oriented measure of cost in the presence of agent diversity; 2) A full characterization of those network topologies for which diversity always helps, for all latency functions and demands.

1 Introduction

It is a common belief that diversity helps. In non-cooperative multi-agent systems, where a central theme is the tension between selfish behavior and social optimality—can diversity of agent preferences help to bring us closer to the coveted social optimality? We provide an answer to this question in the context of non-atomic selfish routing, where diversity naturally arises in the way agents trade-off two criteria, for example, time and cost, or, in the presence of uncertain delays, expectation and variance of delay.

As we shall see, there is no unique answer. Rather, it depends on the setting. To address our question we need to specify how to measure the cost of an outcome, and define a benchmark setting with no agent heterogeneity. To measure the cost of an outcome, we treat an agent’s cost as the sum of two terms associated with two criteria: If we let ℓ_P denote the cost of one criterion (e.g., the latency) over a path P , and v_P be the cost of the second criterion, then the overall cost is given by $\ell_P + r \cdot v_P$, where r is our diversity parameter. The special case of $r = 0$ corresponds to indifference to the second criterion and results in the classic selfish routing model where agents simply minimize travel time.

A first approach to measuring the effect of diversity might be to compare the cost of an outcome with $r = 0$ (i.e., just the total latency) to that with other values of r , including possibly mixed values of r across the population being routed. However, this approach does not pinpoint the gains and losses from agent heterogeneity as opposed to agent homogeneity; rather, it (mostly) pinpoints the gains and losses depending on whether players are affected by the second criterion or not. Instead, we focus on the sum of the costs incurred by the agents as measured by their cost functions, and compare costs incurred by a heterogeneous population of agents to those incurred by an equivalent population of homogeneous agents. What are equivalent populations? Suppose the heterogeneous population’s diversity profile is given by a population density function $f(r)$. Then, we define the corresponding homogeneous population to have the single diversity parameter $\bar{r} = \int r f(r) dr$. In addition, we require the two populations to have the same size, in the sense that the total source-to-sink flows that they induce are equal.

Contribution. We fully characterize the topology of networks for which diversity is never harmful, regardless of the demand size and the distribution of the diversity parameter (discrete or continuous). We do so both for single and multi-commodity networks.

For single-commodity networks it turns out that this topology is that of series-parallel networks. In Theorem 1, we show that if the network is series-parallel, then diversity only helps for any choice of demand and edge functions. The key observation is that there is a path for which the homogeneous flow is at least as large as the heterogeneous flow. As the cost of the homogeneous flow is the same on all used paths, while the cost of each unit of heterogeneous flow is lowest on the path it uses, one can then deduce that the cost of the heterogeneous flow is at most that of the homogeneous flow. To show necessity, we first provide an instance on the Braess graph for which diversity is harmful, and then show how to embed it in any non-series-parallel graph.

In multi-commodity networks, by the result above, each commodity must route its flow through a series-parallel sub-network. But, as Proposition 2 shows, this is not enough, and the way in which these series-parallel networks overlap needs to be constrained. The necessary constraint is exactly captured by the class of *block-matching* networks, defined in this paper. Sufficiency in this case then follows quite easily

from the same result for the single commodity case.

The main technical challenge is to show necessity. To this end, assuming diversity does no harm, we show, via a case analysis, how the subnetworks of the commodities may overlap. First, in Proposition 2, we give an instance on a network of two commodities and three paths for which diversity hurts. Then we mimic this instance on a general network. The difficult part is to choose the corresponding paths for the mimicking, so that, in the created instance, all the flow under both equilibria goes through these paths. The challenge is that the commodities' subnetworks may overlap in subtle ways.

Related work. Our work was inspired by the selfish routing literature where agents have to tradeoff two criteria such as time versus tolls, or expected travel time versus variance. In the former setting, early results (e.g., [Beckmann *et al.*, 1956]) showed that tolls can help implement the social optimum as an equilibrium, when agents all have the same linear objective function combining time and money. Much more recently, these results were extended to the case where agents trade-off travel time and money differently, by Cole *et al.* (2003) and Fleischer (2005) for the single commodity case, and by Karakostas and Kolliopoulos (2004a) and Fleischer *et al.* (2004) for the multicommodity case. We remark that the above works, apart from [Cole *et al.*, 2003], consider only the case where the social welfare is defined as the total travel time, whereas in our work we consider the total agent cost, which encapsulates both criteria. This, for example, is also the case for Christodoulou *et al.* (2014) and Karakostas and Kolliopoulos (2004b). In stochastic selfish routing and related frameworks, most models assume homogeneous agents (e.g., [Piliouras *et al.*, 2013], [Nikolova and Stier-Moses, 2015], [Lianeas *et al.*, 2016], [Kleer and Schäfer, 2016]) and study the degradation of a network's performance due to risk aversion. Fotakis *et al.* (2015) considered games with heterogeneous risk-averse players and showed how uncertainty can be used to improve a network's performance.

Characterizing the topology of networks that satisfy some property is a common theme in computer science. Relevant to our work, Epstein *et al.* (2009) characterized the topology of single-commodity networks for which all Nash equilibria are social optima (under bottleneck costs), and Milchtaich (2006) characterized the topology of single-commodity networks which do not suffer from the Braess Paradox for any cost functions. Chen *et al.* (2015) fully characterized the topology of multi-commodity networks that do not suffer from the Braess paradox. These characterizations appear similar to ours, although there does not seem to be any other connection between the two problems, as (i) there are instances where diversity helps while the Braess Paradox occurs and others where diversity hurts but the paradox does not occur, and (ii) the Braess Paradox may occur in series-parallel networks when considering selfish routing with heterogeneity in agent preferences, which is not the case for the classic selfish routing model.

In the following works, the characterizing topology for the corresponding question (for a single commodity) is similar to ours. Fotakis and Spirakis (2008) considered atomic games and proved that series-parallel networks are the largest class of networks for which strongly optimal tolls are known to

exist. Nikolova and Stier-Moses (2015) considered homogeneous agents and a social cost function that does not account for the second criterion; They showed that series-parallel networks admit the best bound on the degradation of the network due to risk aversion. Theorem 4 of Acemoglu *et al.* (2016), proves that series-parallel networks are the characterizing topology for what they call the Informational Braess Paradox with Restricted Information Sets; this theorem compares the cost of one agent type before and after more information is revealed to agents of that type, but does not consider the change in the cost of other agent types. In contrast, our work considers non-atomic games with heterogeneous agents and bounds the overall costs faced by the collection of agents. Most relevant to our work is [Meir and Parkes, 2014] and its Theorem 3.1 as it implies that for series-parallel networks the cost of an agent of average parameter only increases when switching from the heterogeneous instance to the corresponding homogeneous one and thus for our sufficiency theorems, one is left to prove that the heterogeneous equilibrium cost is no greater than the cost of an agent of average parameter (though we give a different proof).

2 Preliminaries

Consider a directed multi-commodity network $G = (V, E)$ with an aggregate demand of d_k units of flow between origin-destination pairs (s_k, t_k) for $k \in K$. We let \mathcal{P}_k be the set of all paths between s_k and t_k , and $\mathcal{P} := \cup_{k \in K} \mathcal{P}_k$ be the set of all origin-destination paths. We let $[m]$ denote $\{1, \dots, m\}$. We assume that $K = [m]$, for some m . The agents in the network—i.e., the players of the game—must choose routes that connect their origins to their destinations. We encode the collective decisions of agents in a flow vector $f = (f_\pi)_{\pi \in \mathcal{P}} \in \mathbb{R}_+^{|\mathcal{P}|}$ over all paths. Such a flow is feasible when demands are satisfied, as given by constraints $\sum_{\pi \in \mathcal{P}_k} f_\pi = d_k$ for all $k \in K$. For simplicity, we let f_e denote the flow on edge e ; note that $f_e = \sum_{\pi: e \in \pi} f_\pi$. When we need multiple flow variables, we use the analogous notation g, g_π, g_e .

The network is subject to congestion that affects two criteria the players consider. These two criteria are modeled by two *edge-dependent functions* that take as input the edge flow f_e of e , for each edge e : a latency function $\ell_e(x)$ assumed to be continuous and non-decreasing, and a deviation function $\sigma_e(x)$ assumed to be continuous (but not necessarily non-decreasing).

Throughout the paper we refer to the agent's objective as the cost along a route. Formally, for a given agent, on letting $\ell_\pi(f) = \sum_{e \in \pi} \ell_e(f_e)$ and $\sigma_\pi(f) = \sum_{e \in \pi} \sigma_e(f_e)$, for a constant $r \geq 0$ that quantifies the agent diversity parameter, the agent's cost along route π under flow f is

$$c_\pi^r(f) = \sum_{e \in \pi} \ell_e(f_e) + r \sum_{e \in \pi} \sigma_e(f_e) = \ell_\pi(f) + r\sigma_\pi(f) \quad (1)$$

We assume that for any edge and any agent's diversity parameter r , the functions ℓ_e and $\ell_e + r\sigma_e$ are non-decreasing. Note that if there is an upper bound r_{\max} on the possible values of the diversity parameter r , then the latter assumptions do not require σ_e to be non-decreasing. This is desirable because, for example, in risk-averse selfish routing where σ_e models the variance, σ_e may be decreasing in the flow.

Players Heterogeneity. We assume that there may be more than one value of the diversity parameters r for the players routing commodity k . We use the term *single-minded* to refer to players with $r = 0$. We handle both the cases of a continuous and a discrete distribution of the diversity parameter among the players, though, for our results, we need only consider the discrete case.¹

For a discrete distribution of, say, n discrete values r_1^k, \dots, r_n^k , the demand d_k is a vector $d_k = (d_1^k, \dots, d_n^k)$ where each d_i^k denotes the total demand of Commodity k with diversity parameter r_i^k . We let d^k denote Commodity k 's total demand, $d^k = \sum_{i=1}^n d_i^k$. Variables f_π^r and f_e^r denote the flow of diversity parameter r on path π and edge e , respectively. Formally, an instance is described by the tuple $(G, \{(\ell_e, \sigma_e)\}_{e \in E}, \{(s_k, t_k)\}_{k \in K}, \{d_k\}_{k \in K}, \{r_k\}_{k \in K})$, where $r_k = (r_1^k, \dots, r_n^k)$ is the vector of different diversity parameters encountered in the heterogeneous population.

Equilibrium flows. The Wardrop equilibrium of an instance is a flow f such that for every $k \in K$, for every path $\pi \in \mathcal{P}_k$ with positive flow, and any diversity parameter r on it, the path cost $c_\pi^r(f) \leq c_{\pi'}^r(f)$ for all paths $\pi' \in \mathcal{P}_k$.

From here on, we shall refer to the Wardrop equilibrium as the equilibrium. Our goal is to compare the total agent cost at the equilibrium of an instance with a population that has heterogeneous diversity parameters, to the total agent cost at the equilibrium of the same instance but with the population of each commodity keeping its magnitude but changed to be homogeneous, with diversity parameter equal to the expected value of the diversity parameter distribution in the heterogeneous population of the commodity. To differentiate more easily, for a heterogeneous instance we call the former the *heterogeneous equilibrium* and the latter the (corresponding) *homogeneous equilibrium*. We usually denote the heterogeneous equilibrium by g and the homogeneous equilibrium by f . The existence of both equilibria is guaranteed by e.g. [Schmeidler, 1973, Theorem 2]. We note here that in general we do not need uniqueness of equilibria, neither for the edge costs nor for the edge flows. Our results hold for any arbitrary pair of heterogeneous and homogeneous equilibria of the corresponding instances. Also, as for classic routing games, without loss of generality (WLOG) we may assume that equilibrium flows are acyclic.

Total Costs. For a heterogeneous equilibrium flow vector g , the *heterogeneous total cost* of Commodity k is denoted by $C^{k,ht}(g) = \sum_{j=1}^n d_j^k c_j^{k,r_j^k}(g)$ where $c_j^{k,r_j^k}(g)$ denotes the common cost at equilibrium g for players of diversity parameter r_j^k in Commodity k . The *heterogeneous total cost* of g is then $C^{ht}(g) = \sum_{k \in K} C^{k,ht}(g)$. For the corresponding homogeneous equilibrium flow f , i.e. the instance with diversity parameter \bar{r}^k , where \bar{r}^k denotes the average diversity parameter for Commodity k , players of Commodity k share the same cost $c^{\bar{r}^k}(f)$. Then, the *homogeneous total cost* of Commodity k under f is $C^{k,hm}(f) = d_k c^{\bar{r}^k}(f)$, and the *homogeneous total cost* of f is $C^{hm}(f) = \sum_{k \in K} C^{k,hm}(f)$.

¹This is basically because the flow g_p of a path p with diversity parameter within some range, can be changed to have parameter equal to the average parameter on p , and the equilibrium remains.




Figure 1: A block-matching network of 2 commodities. G_1 and G_2 are series-parallel and their block representations are $G_1 = s_1 A u B v C w D t_1$ and $G_2 = s_2 E w D t_1 F u B v G t_2$. G_1 and G_2 share exactly blocks B and D and do not share any edge on any other of their blocks. If we add an edge from s_1 to t_1 , then the network stops being block-matching since G_1 will be a block by itself and it will not match any of the blocks of G_2 .

Finally, if $C^{ht}(g) \leq C^{hm}(f)$, we say that *diversity helps*; if not, we say that *diversity hurts*. For our characterization to be meaningful, we assume an *average-respecting demand*, i.e., a demand where $\forall i, j : \bar{r}^i = \bar{r}^j$. Otherwise, diversity may hurt in simple instances, e.g., with two parallel links and two commodities.

Networks. For a network G we let $V(G)$ and $E(G)$ denote its vertex set and edge set, respectively.

A directed $s-t$ network G is *series-parallel* if it consists of a single edge (s, t) , or it is formed by the series or parallel composition of two series-parallel networks with terminals (s_1, t_1) and (s_2, t_2) , respectively. In a *series composition*, t_1 is identified with s_2 , s_1 becomes s , and t_2 becomes t . In a *parallel composition*, s_1 is identified with s_2 and becomes s , and t_1 is identified with t_2 and becomes t . The internal vertices of a series-parallel network G are all its vertices other than its terminals.

An $s-t$ series-parallel network may be represented using a sequence of networks B_j connected in series, where each B_j is either a single edge or two series-parallel networks connected in parallel. Given a series-parallel network H , we can write $H = s B_1 v_1 B_2 v_2 \dots B_{b-1} v_{b-1} B_b t$, where for any j and triple $x B_j y$, x and y are the terminals of the series-parallel network B_j , and B_j is either a single edge or a parallel combination of two series-parallel networks. We refer to the B_j 's as *blocks*, the prescribed representation as the *block representation* of H , and the v_i 's as *separators*, as they separate s from t . Two series-parallel networks G_1 and G_2 are said to be *block-matching* if for every block B of G_1 and every block D of G_2 , either $E(B) = E(D)$ or $E(B) \cap E(D) = \emptyset$. Note that $E(B) = E(D)$ implies that B and D have the same terminals and direction, as for either B or D , the source vertex will have only outgoing edges toward the internal vertices and the target vertex will have only incoming edges from the internal vertices.

For a k -commodity network G , let G_i be the subnetwork of G that contains all the vertices and edges of G that belong to a simple s_i-t_i path for Commodity i . In other words, G_i is the subnetwork of G for Commodity i that equilibria flows will consider, as they are, WLOG, acyclic. A multi-commodity network G is *block-matching* if for every i , G_i is series-parallel, and for every i, j , G_i and G_j are block-matching. An example is given in Figure 1.

3 Topology of Single-Commodity Networks for which Diversity Helps

In this section, we fully characterize the topology of single-commodity networks for which, with any choice of heterogeneous demand and edge functions, diversity helps. WLOG we may restrict our attention to single-commodity networks whose edges all belong to some simple source-destination path as only these edges are going to be used by the (WLOG, acyclic) equilibria and thus all other edges can be discarded. It turns out that this topology is exactly that of series-parallel networks (Theorems 1 and 2).

3.1 Series Parallel Implies Diversity is Helpful

Throughout this section we will be considering a heterogeneous instance \mathcal{G} on an s - t series-parallel network G . We let \mathcal{F} denote the corresponding homogeneous instance. We let g denote an equilibrium flow for \mathcal{G} and f an equilibrium flow for \mathcal{F} . Finally, we let $C^{ht}(g)$ denote the cost of flow g and $C^{hm}(f)$ the cost of flow f . Although redundant, we keep the superscripts as a further reminder of the flow type at hand.

The key observation is that there is a path P used by flow f such that for every edge in P , $f_e \geq g_e$, and hence for any $r \in [0, r_{\max}]$, $c_p^r(f) \geq c_p^r(g)$ (Lemmas 1² and 2). We then deduce our result: $C^{ht}(g) \leq C^{hm}(f)$ (Theorem 1).

Lemma 1. *Let G be an s - t series-parallel network and let x and y be flows on G that route d_1 and d_2 units of traffic respectively, with $d_1 \geq d_2$ and $d_1 > 0$. Then, there exists an s - t path P such that for all $e \in P$, $x_e > 0$ and $x_e \geq y_e$.*

Lemma 2. *There exists a path P used by f such that for any $r \in [0, r_{\max}]$, $c_p^r(g) \leq c_p^r(f)$.*

Proof. Flows f and g have the same magnitude on the series-parallel network G . Applying Lemma 1 with $x = f$ and $y = g$ implies that there exists an s - t path P such that for all $e \in P$, $f_e > 0$, implying that WLOG P is used by f , and $f_e \geq g_e$. By assumption, for any $r \in [0, r_{\max}]$, $\ell_e + r\sigma_e$ is non-decreasing, and thus for all $e \in P$, $\ell_e(f_e) + r\sigma_e(f_e) \geq \ell_e(g_e) + r\sigma_e(g_e)$. Consequently, $\sum_{e \in P} (\ell_e(f_e) + r\sigma_e(f_e)) \geq \sum_{e \in P} (\ell_e(g_e) + r\sigma_e(g_e)) \Leftrightarrow c_p^r(g) \leq c_p^r(f)$ as needed. \square

Theorem 1. $C^{ht}(g) \leq C^{hm}(f)$.³

Proof. Since G is a series-parallel network, on setting $\bar{r} = E[r]$ and then applying Lemma 2, we obtain that there is a path P used by f such that

$$\ell_p(f) + \bar{r}\sigma_p(f) \geq \ell_p(g) + \bar{r}\sigma_p(g) \quad (2)$$

WLOG we can assume that the total demand $d = 1$. We first bound the total cost of g in terms of the cost of path P under g and then we use (2) to further bound it in terms of the cost of path P under f . The latter equals the cost of f , as the demand is equal to 1.

²Lemma 1 is similar to [Milchtaich, 2006, Lemma 2].

³The inequality might be strict. Consider the case of 2 parallel links with $(\ell_1(x), \sigma_1(x)) = (1, x)$ and $(\ell_2(x), \sigma_2(x)) = (2, 0)$, and 1 unit of flow, half with $r = 0$ and half with $r = 2$.




Figure 2: The Braess network of Proposition 1.

Consider the heterogeneous equilibrium flow g . By the equilibrium conditions, for any player of diversity parameter r , for any r , the cost she incurs with flow g is $c^r(g) \leq \sum_{e \in P} \ell_e(g_e) + r \sum_{e \in P} \sigma_e(g_e)$. In other words, there is no incentive to deviate to path P (if not already on it). Thus, for the demand vector (d_1, \dots, d_k) of diversity parameters (r_1, \dots, r_k) , $C^{ht}(g) \leq \sum_{i=1}^k d_i (\sum_{e \in P} \ell_e(g_e) + r_i \sum_{e \in P} \sigma_e(g_e)) = \ell_p(g) + \bar{r}\sigma_p(g)$, with the equality following as the total demand is 1 and the average diversity parameter is $\bar{r} = \sum_{i=1}^k d_i r_i$. As P is used by f , we have $C^{hm}(f) = \ell_p(f) + \bar{r}\sigma_p(f)$, and applying (2) we obtain $C^{ht}(g) \leq \ell_p(g) + \bar{r}\sigma_p(g) \leq \ell_p(f) + \bar{r}\sigma_p(f) = C^{hm}(f)$. \square

3.2 The Series Parallel Condition is Necessary

To prove the necessity of the network being series-parallel, we begin by constructing an instance for which diversity hurts, i.e. the heterogeneous equilibrium has total cost strictly greater than the total cost of the homogeneous equilibrium (Proposition 1). Then, in Theorem 2, we show how to embed this instance into any network that is not series-parallel.

Recall the Braess graph G_B , shown in Figure 2.

Proposition 1. *For any strictly heterogeneous demand on the Braess graph G_B , there exist edge functions $(\ell_e)_{e \in E}$ and $(\sigma_e)_{e \in E}$ that depend on the demand, for which $C^{ht}(g) > C^{hm}(f)$.⁴*

Proof. We may assume WLOG that the demand is of unit size. Let \bar{r} be the average diversity parameter, r_0 be the minimum of the diversity parameters' distribution and let d_0 be the total demand with diversity parameter equal to r_0 . Since the demand is strictly heterogeneous, it must be $r_0 < \bar{r}$. In addition, we let h be any continuous, strictly increasing cost function with $h(\frac{1}{2}) = 1$ and $h(\frac{1}{2} + \frac{d_0}{2}) = 1 + \frac{\bar{r} - r_0}{2}$.

Consider the Braess graph $G_B = (\{s, u, v, t\}, \{(s, u), (u, t), (u, v), (s, v), (v, t)\})$ with cost functions $\ell_{(s,u)}(x) = \ell_{(v,t)}(x) = h(x)$, $\sigma_{(s,u)}(x) = \sigma_{(v,t)}(x) = 0$, $\ell_{(u,t)}(x) = \ell_{(s,v)}(x) = 2 + \frac{\bar{r} + r_0}{2}$, and $\sigma_{(u,t)}(x) = \sigma_{(s,v)}(x) = 0$, and $\ell_{(u,v)}(x) = 1$ and $\sigma_{(u,v)}(x) = 1$. The instance is shown in Figure 2.

The heterogeneous equilibrium g routes the d_0 units that have $r = r_0$ through the zig-zag path, i.e. path s, u, v, t ; the rest of the flow is split between the upper and lower paths s, u, t and s, v, t , and has $C^{ht}(g) = 3 + \bar{r}$. The homogeneous equilibrium f splits the flow between the upper and lower paths and has $C^{hm}(f) = 3 + \frac{\bar{r} + r_0}{2} < 3 + \bar{r} = C^{ht}(g)$. \square

⁴The proposition remains true even if we are restricted to only using affine functions.

Theorem 2. *If G is not series-parallel, then for any strictly heterogeneous demand there are cost functions for which $C^{ht}(g) > C^{hm}(f)$.*

Proof sketch. If G is not series-parallel then the Braess graph can be embedded in it (see e.g. [Valdes *et al.*, 1979]). Thus, starting from the Braess network G_B , by subdividing edges, adding edges and extending one of the terminals by one edge, we can obtain G . Fixing such a sequence of operations, for the given heterogeneous demand, we start from the Braess instance given by Proposition 1 and apply the sequence of operations one by one, so that, throughout the procedure, there are exactly 3 paths that correspond to the upper, lower and zig-zag paths of Proposition 1, with corresponding costs, and all other paths have some very large costs. Then, $C^{ht}(g) > C^{hm}(f)$ follows exactly as in Proposition 1. \square

4 Topology of Multi-Commodity Networks for which Diversity Helps

In this section we fully characterize the topology of multi-commodity networks for which, with any choice of heterogeneous average-respecting demand and edge functions, diversity helps. Because of Theorem 2, if we require diversity to help on any instance on G , then for any commodity i , G_i needs to be series-parallel. Yet, as we shall see in Proposition 2, this is not enough. We also need to understand the overlaps of the G_i 's. It turns out that the allowable overlaps are exactly captured by the topology of block-matching networks (Theorems 3 and 4).

4.1 Sufficiency

Using Theorem 1, we can obtain an analogous theorem for the multi-commodity case.

Theorem 3. *Let G be a k -commodity block-matching network. Then, for any instance on G with average-respecting demand $C^{ht}(g) \leq C^{hm}(f)$.*

Proof sketch. Consider Commodity i and let $G_i = s_i B_1 v_1 \dots v_{b_i-1} B_{b_i} t_i$ be its block representation. Consider an arbitrary B_j with terminals v_{j-1} and v_j . Because G is block-matching, any other Commodity l either contains B_j as a block in its block representation or contains none of its edges. Also, recall that, as explained in the preliminaries section, if G_l contains B_j , it has the same terminals v_{j-1} and v_j . This implies that under any routing of the demand, either all of l 's demand goes through B_j or none of it does. This means that under both equilibria g and f , the total traffic routed from v_{j-1} to v_j through B_j is the same which further implies that, if restricted to the block, the cost of the heterogeneous equilibrium is less than or equal to that of the homogeneous equilibrium: $C^{ht}(g)|_{B_j} \leq C^{hm}(f)|_{B_j}$. On the other hand, if we let \mathcal{B} be the set of all the blocks of all commodities, then $C^{ht}(g) = \sum_{B \in \mathcal{B}} C^{ht}(g)|_B$ and $C^{hm}(f) = \sum_{B \in \mathcal{B}} C^{hm}(f)|_B$ which using the previous inequality proves the result. \square




Figure 3: The network for Proposition 2

4.2 Necessity

To derive the necessity we first give an example of a non-block-matching network for which diversity hurts (Proposition 2). Then, after giving some properties for commodities for which the corresponding G_i 's are series-parallel (Lemmas 3 and 4), we mimic the above example to obtain contradicting instances for networks that are not block-matching and thereby prove Theorem 4.

Let G be the following 2-commodity network, depicted in Figure 3. G_2 , the subnetwork for Commodity 2, consists of a simple s_2-t_2 path P_2 , while G_1 , the subnetwork for Commodity 1, is formed from two simple s_1-t_1 paths named P_1 and P_3 ; P_1 and P_2 are disjoint, while P_2 and P_3 share a single edge, named e_2 . Finally e_1 is an edge on P_1 but not on P_3 .

Proposition 2. *There exist edge functions and demands on G for which diversity hurts.*

Proof. Let $d_1 = d_2 = 1$ be the total demands for Commodity 1 and 2 respectively. Let G_1 's demand consist of $\frac{3}{4}$ single-minded players (i.e., with diversity parameter equal to 0) and $\frac{1}{4}$ players with diversity parameter equal to 4, and let G_2 's demand be homogeneous with diversity parameter equal to 1. To all edges other than e_1 and e_2 , assign latency and deviation functions equal to 0. Assign edge e_1 the constant latency function $\ell_1(x) = 1$ and the constant deviation function $\sigma_1(x) = 2$. Assign edge e_2 the constant deviation $\sigma_2 = 0$, and as latency function any ℓ_2 that is continuous and strictly increasing, with $\ell_2(1) = 3$ and $\ell_2(\frac{5}{4}) = 9$.

The equilibrium costs depend only on the flow through edges e_1 and e_2 , as all other edges have cost 0. Also note that at least 1 unit of flow will go through e_2 as this is the only route for G_2 's demand.

In the heterogeneous equilibrium g of this instance, $\frac{3}{4}$ units of flow are routed through e_1 , and $1 + \frac{1}{4}$ units of flow are routed through e_2 , giving $C^{ht}(g) = 1 \cdot \frac{3}{4}d_1 + 9 \cdot \frac{1}{4}d_1 + 9 \cdot d_2 = 12$. In the homogeneous equilibrium f , G_1 's demand uses only P_1 . Thus, $C^{hm}(f) = 3 \cdot d_1 + 3 \cdot d_2 = 6 < C^{ht}(g)$. \square

Remark 1. *Proposition 2 would still hold if the common portion of P_1 and P_3 , and the portion of P_2 after e_2 , both had positive costs instead of zero. This is close to the way we will mimic this instance in the proof of Theorem 4. The idea, in both equilibria, is to route all the flow of Commodity 1 through two paths, P_1 and P_3 , each containing one of e_1 or e_2 , and to route the flow of Commodity 2 through a path, P_2 , that contains e_2 . This is done by putting (relatively) big constants as latency functions on all other edges that depart from vertices of the corresponding paths up to the point where e_1 or e_2 is reached, though some care is needed. Then, the relation of the equilibria costs will follow as in Proposition 2, as the exact same edge functions will be used for edges e_1*

and e_2 . This will be specified precisely when we give the construction.

Next, we state some useful properties of series-parallel networks that are based on their block structure. They will be used in the proof of Theorem 4.

Lemma 3. *Let i be a commodity of network G and suppose that G_i is series-parallel.*

(i) *Let B_1 and B_2 be distinct blocks of G_i , with B_1 preceding B_2 . There is no edge in G from an internal vertex of B_1 to an internal vertex of B_2 .*

(ii) *Let u and v be vertices in G_i . If (u, v) is an edge of G then there is a simple s_i - t_i path in G_i that contains both u and v (not necessarily in that order).*

Lemma 4. *Let i be a commodity of network G and suppose that G_i is series-parallel with block representation $G_i = s_i B_1 v_1 \dots v_{b-1} B_b t_i$. Let w be a vertex of B_k for some $k \in [b]$.*

(i) *Suppose that $w \neq v_{k-1}$, and let P be an arbitrary path from a vertex u , in a block A that precedes B_k in the block representation, to vertex w . Let w' be the first vertex on P that is an internal vertex in B_k , if any. Then P must include an edge of B_k exiting v_{k-1} prior to visiting w' .*

(ii) *Suppose that $w \neq v_k$. Then any path of G from w to a vertex u in a block succeeding B_k has to first enter v_k through one of its incoming edges that belong to B_k , before going to a block C that succeeds B_k in the block representation.*

(iii) *Every simple v_{k-1} - v_k path in G is completely contained in B_k .*

Theorem 4. *Let G be a multi-commodity network. If diversity helps for every instance on G with average-respecting demand (i.e. for any heterogeneous equilibrium g and any homogeneous equilibrium f , $C^{ht}(g) \leq C^{hm}(f)$), then G is a block-matching network.*

Proof. Let G have k commodities. First, we note that for any $i \in [k]$, G_i is a series-parallel network. For otherwise, by Proposition 1, there is some heterogeneous players' demand for Commodity i and edge functions for G_i such that diversity hurts. Letting all other commodities have zero demand yields an instance on G for which diversity hurts, a contradiction.

To prove that G is block-matching, it remains to show that for any two commodities i and j of G , for any block B of G_i and any block D of G_j , either $E(B) = E(D)$ or $E(B) \cap E(D) = \emptyset$. To reach a contradiction we assume otherwise, i.e. WLOG we assume that for Commodities 1 and 2 there exist two blocks B of G_1 and D of G_2 that share some common edge, and at the same time, WLOG, there is an edge in B that is not in D . The latter implies that B is not a single edge, and thus it must be a parallel combination of two series-parallel networks.

Let u and v be the endpoints of B . We first prove that all simple s_2 - t_2 paths of G_2 that share an edge with B first traverse an edge starting at u before traversing any other edge of B (Proposition 3). Then we prove that all s_2 - t_2 simple paths of G_2 , that share an edge with B , reach u before traversing any internal vertex of B (Proposition 4). Since $E(B) \cap E(D) \neq \emptyset$, there is a simple s_2 - t_2 path of G_2 that shares an edge with B . Proposition 4 implies that this path, Q ,




Figure 4: Illustrating why $P_2 \neq P$ in general in Proposition 3

has a subpath consisting of a simple s_2 - u path Q_1 that shares no internal vertex with B . A completely symmetric argument shows that Q has a subpath consisting of a simple v - t_2 path Q_3 that shares no internal vertex with B .⁵ But then, for any simple u - v path Q_2 inside B , the path $Q' = Q_1, Q_2, Q_3$ is a simple s_2 - t_2 path, and thus it belongs to G_2 . But this implies that all the edges of B belong to G_2 and because B is a block, these edges will all be in a single block of G_2 . This block must be block D , since by assumption $E(B) \cap E(D) \neq \emptyset$, contradicting the existence of an edge in B and not in D . Therefore, once these propositions are proved, the theorem will follow.

The proofs of these propositions rely on the same idea. For each proposition, assuming that it does not hold, we construct instances, i.e. we choose demand and edge functions for G , such that diversity hurts, contradicting the assumption that for any instance on G diversity helps. The construction of the contradicting instances is based on Remark 1.

Proposition 3. *Let P be a simple s_2 - t_2 path in G_2 which shares an edge with B . The first edge on P in B departs from u , i.e. has the form (u, x) for some x in B .*

Proof. Let B be the parallel combination of H_1 and H_2 . WLOG we may assume that P only visits vertices of G_1 , plus s_2 and t_2 , as we may treat subpaths of P that have vertices that lie outside G_1 as simple edges. Let w be the first internal vertex of P that belongs to B , and WLOG suppose that w lies in H_1 . By Lemma 3(ii), the edge of P exiting w will either go toward t_1 , i.e. forward, and thus traverse an edge of B for the first time (recall also Lemma 4(ii)), or will go toward s_1 , i.e. backward, either staying in H_1 or going back to one of the preceding blocks of B . If it goes to one of the preceding blocks of B , then by Lemma 4(i), it has to traverse an edge of B departing from u in order to re-enter the internal portion of B (recall that P has some edge in B) and then the proposition would hold. The remaining possibility is that the backward edge leads to another internal vertex of H_1 . However, we can only repeat this process finitely often so if the proposition does not hold, it must be that P eventually traverses a first edge in B that departs from an internal vertex of H_1 . In this case we will reach a contradiction by creating an instance where diversity hurts. This instance will be based on the instance of Proposition 2.

We would like to use the following construction at this point. Let P_2 be the path P resulting from the discussion in

⁵For the symmetric argument, simply reverse all the arcs and the directions of the demand.

the previous paragraph and let e_2 be the first edge on P that lies in B . Then let P_3 be an s_1-t_1 path through e_2 . Recall that e_2 lies in H_1 . Now let P_1 be an s_1-t_1 path that goes through H_2 and let e_1 be an arbitrary edge on P_1 in B . The intention is to force the s_1-t_1 flow to use just paths P_1 and P_3 , while the s_2-t_2 flow uses just path P_2 , at the same time ensuring that diversity is harmful as in Proposition 2. Consider the following edge functions. e_1 and e_2 receive the same edge functions as in Proposition 2. The other edges all receive a constant deviation function equal to 0. For their latency functions, edges on P_1 and P_3 that are in B receive 0 functions. Out-edges from P_1 and P_3 that lie in B all receive functions of constant value N even if they are on P_2 . All as yet unassigned edges on P_2 receive 0 functions, and the remaining edges are all given functions of constant value $M \gg N$. However, the example in Figure 4 shows that there is a zero cost s_2-t_2 path (s_2, w, z, t_2) , which defeats the construction.

We fix this problem by defining the path P_2 as follows. Let x be the first vertex on path P (in the example, this is w) such that there is an edge (x, y) in B and such that there is a $y-t_2$ path P_{y-t_2} which does not go through any earlier vertex on P (i.e. any vertex from s_2 to x inclusive). Then (x, y) is chosen to be e_2 , and P_2 is defined to be the simple path comprising the initial portion of P up to x , followed by e_2 , followed by P_{y-t_2} (it may be that $P_2 = P$). Now the above cost functions, modulo a few details, achieve the desired contradiction. \square

Proposition 4. *All simple s_2-t_2 paths of G_2 that share an edge with B reach u before any internal vertex of B .*

The proof is largely similar to that of Proposition 3. \square

Acknowledgements

We thank the anonymous referees for their thoughtful feedback that helped improve this work. Part of this work was completed while the authors were visiting the Simons Institute for the Theory of Computing, Berkeley, CA. Research partially supported by NSF grants CCF-1527568, CCF-1216103, CCF-1350823, CCF-1331863, CCF 1733832.

References

- [Acemoglu *et al.*, 2016] Daron Acemoglu, Ali Makhdomi, Azarakhsh Malekian, and Asuman E. Ozdaglar. Informational braess’ paradox: The effect of information on traffic congestion. *CoRR*, abs/1601.02039, 2016.
- [Beckmann *et al.*, 1956] Martin Beckmann, C. B. McGuire, and Christopher B. Winsten. *Studies in the economics of transportation*. Yale University Press, 1956.
- [Chen *et al.*, 2015] Xujin Chen, Zhuo Diao, and Xiao-Dong Hu. Excluding braess’s paradox in nonatomic selfish routing. In *SAGT ’15*, 2015.
- [Christodoulou *et al.*, 2014] George Christodoulou, Kurt Mehlhorn, and Evangelia Pyrga. Improving the price of anarchy for selfish routing via coordination mechanisms. *Algorithmica*, 2014.
- [Cole *et al.*, 2003] Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. Pricing network edges for heterogeneous selfish users. In *STOC ’03*, 2003.
- [Epstein *et al.*, 2009] Amir Epstein, Michal Feldman, and Yishay Mansour. Efficient graph topologies in network routing games. *Games and Economic Behavior*, 66, 2009.
- [Fleischer *et al.*, 2004] Lisa Fleischer, Kamal Jain, and Mohammad Mahdian. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. *FOCS ’04*, 2004.
- [Fleischer, 2005] Lisa Fleischer. Linear tolls suffice: New bounds and algorithms for tolls in single source networks. *Theoretical Computer Science*, 2005.
- [Fotakis and Spirakis, 2008] Dimitris Fotakis and Paul G. Spirakis. Cost-balancing tolls for atomic network congestion games. *Internet Mathematics*, 5(4):343–363, 2008.
- [Fotakis *et al.*, 2015] Dimitris Fotakis, Dimitris Kalimeris, and Thanasis Lianeas. Improving selfish routing for risk-averse players. In *WINE ’15*, 2015.
- [Karakostas and Kolliopoulos, 2004a] George Karakostas and Stavros G. Kolliopoulos. Edge pricing of multicommodity networks for heterogeneous selfish users. *FOCS ’04*, 2004.
- [Karakostas and Kolliopoulos, 2004b] George Karakostas and Stavros G. Kolliopoulos. The efficiency of optimal taxes. In *CAAN ’04*, 2004.
- [Kleer and Schäfer, 2016] Pieter Kleer and Guido Schäfer. The impact of worst-case deviations in non-atomic network routing games. In *SAGT ’16*, pages 129–140, 2016.
- [Lianeas *et al.*, 2016] Thanasis Lianeas, Evdokia Nikolova, and Nicolas E. Stier-Moses. Asymptotically tight bounds for inefficiency in risk-averse selfish routing. In *IJCAI ’16*, 2016.
- [Meir and Parkes, 2014] Reshef Meir and David C. Parkes. Playing the wrong game: Bounding externalities in diverse populations of agents. *CoRR*, abs/1411.1751, 2014. To appear in *AAMAS ’18*.
- [Milchtaich, 2006] Igal Milchtaich. Network topology and the efficiency of equilibrium. *Games and Economic Behavior*, 57, 2006.
- [Nikolova and Stier-Moses, 2015] Evdokia Nikolova and Nicolas E. Stier-Moses. The burden of risk aversion in mean-risk selfish routing. In *EC ’15*, 2015.
- [Piliouras *et al.*, 2013] Georgios Piliouras, Evdokia Nikolova, and Jeff S. Shamma. Risk sensitivity of price of anarchy under uncertainty. In *EC ’13*, 2013.
- [Schmeidler, 1973] David Schmeidler. Equilibrium points of nonatomic games. *Journal of Statistical Physics*, 7, 1973.
- [Valdes *et al.*, 1979] Jacobo Valdes, Robert Endre Tarjan, and Eugene L. Lawler. The recognition of series parallel digraphs. In *STOC ’79*, 1979.