Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Informational Braess' Paradox: The Effect of Information on Traffic Congestion

Daron Acemoglu, Ali Makhdoumi, Azarakhsh Malekian, Asuman Ozdaglar

informs
http://pubsonline.informs.org/journal/opre/

OPERATIONS RESEARCH
Vol. 66, No. 4, July–August 2018, pp. 893–917
ISSN 0030-364X (print), ISSN 1526-5463 (online)

CROSSCUTTING AREAS

Informational Braess' Paradox: The Effect of Information on Traffic Congestion

Daron Acemoglu,[a] Ali Makhdoumi,[b] Azarakhsh Malekian,[c] Asuman Ozdaglar[b]

[a] Department of Economics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139; [b] Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139; [c] Rotman School of Management, University of Toronto, Toronto, Ontario M5S 3E6, Canada

Contact: daron@mit.edu (DA); makhdoum@mit.edu, http://orcid.org/0000-0001-5422-0314 (AlM); azarakhsh.malekian@rotman.utoronto.ca, http://orcid.org/0000-0001-9464-746X (AzM); asuman@mit.edu, http://orcid.org/0000-0002-1827-1285 (AO)

Received: January 8, 2016
Revised: January 18, 2017; August 4, 2017
Accepted: October 5, 2017
Published Online in Articles in Advance: July 23, 2018

Subject Classifications: network/graphs: theory; technology transportation: models, network; games/group decisions: nonatomic; noncooperative
Area of Review: Invited Papers

https://doi.org/10.1287/opre.2017.1712

Abstract. To systematically study the implications of additional information about routes provided to certain users (e.g., via GPS-based route guidance systems), we introduce a new class of congestion games in which users have differing information sets about the available edges and can only use routes consisting of edges in their information set. After defining the notion of an information-constrained wardrop equilibrium (ICWE) for this class of congestion games and studying its basic properties, we turn to our main focus: whether additional information can be harmful (in the sense of generating greater equilibrium costs/delays). We formulate this question in the form of an informational Braess' paradox (IBP), which extends the classic Braess' paradox in traffic equilibria and asks whether users receiving additional information can become worse off. We provide a comprehensive answer to this question showing that in any network in the series of linearly independent (SLI) class, which is a strict subset of series-parallel networks, the IBP cannot occur, and in any network that is not in the SLI class, there exists a configuration of edge-specific cost functions for which the IBP will occur. In the process, we establish several properties of the SLI class of networks, which include the characterization of the complement of the SLI class in terms of embedding a specific set of networks, and also an algorithm that determines whether a graph is SLI in linear time. We further prove that the worst-case inefficiency performance of ICWE is no worse than the standard Wardrop equilibrium.

Funding: This research was supported by the National Science Foundation Cyber-Physical Systems (CPS) Frontiers project "Foundations Of Resilient CybEr-physical Systems (FORCES)" [Award 1239166].

Keywords: Braess' paradox • Wardrop equilibrium • information • traffic • transportation

1. Introduction

The advent of GPS-based route guidance systems, such as Waze or Google maps, promises a better traffic experience to its users, as it can inform them about routes that they were not aware of or help them choose dynamically between routes depending on recent levels of congestion. Although other drivers might plausibly suffer increased congestion as the routes they were using become more congested as a result of this reallocation of traffic, or certain residents may experience elevated noise levels in their side streets, it is generally presumed that the users of these systems (and perhaps society as a whole) will benefit. In this paper, we present a framework for systematically analyzing how changes in the information sets of users in a traffic network (e.g., because of route guidance systems) impact the traffic equilibrium, and we show the conditions under which even those with access to additional information may suffer greater congestion.

Our formal model is a version of the well-known congestion games, augmented with multiple types of users (drivers), each with a different information set about the available edges in the network. These different information sets represent the differing knowledge of drivers about the road network, which may result from their past experiences, from inputs from their social network, or from the different route guidance systems they might rely on. A user can only utilize a route (path between origin and destination) consisting of edges belonging to her information set. Each edge is endowed with a latency/cost function representing costs due to congestion. We generalize the classic notion of the *Wardrop equilibrium* (Wardrop 1952, Beckmann et al. 1956 and Schmeidler 1973), where each (nonatomic) user takes the level of congestion on all edges as given and chooses a route with minimum cost (defined as the summation of costs of edges on the route). Our notion of an *information-constrained Wardrop equilibrium* (*ICWE*), also imposes the same equilibrium condition as Wardrop equilibrium but only for routes that are contained in the information set of each type of user.

After establishing the existence and essential uniqueness of the ICWE and characterizing its main properties for networks with a single origin–destination pair (an assumption we impose for simplicity and later relax), we turn to our key question of whether expanding the information sets of some group of users can make them worse off—in the sense of increasing the level of congestion they suffer in equilibrium. For this purpose, we define the notion of an *informational Braess' paradox* (*IBP*), designating the possibility that users with expanded information sets experience greater equilibrium cost. We then provide a tight characterization of when IBP is and is not possible in a traffic network.

Our main result is that IBP does not occur if and only if the network is *series of linearly independent* (*SLI*). More specifically, this result means that in an SLI network, IBP can never occur, ensuring that users with an expanded information set always benefit from their additional information. Conversely, if the network is not SLI, then there exists a configuration of latency/cost functions for edges for which the IBP will occur. To understand this result, let us consider what the relevant class of networks comprises. The set of SLI networks is a subset of *series-parallel* networks, which are those for which two routes never pass through any edge in opposite directions. An SLI network is obtained by joining together a collection of *linearly independent* (*LI*) networks in a series. LI networks are those in which each route includes at least one edge that is not part of any other route. The intuition of our main result is as follows. To show that the IBP does not occur in an SLI network, we show the result on each of its LI parts. A key property of LI networks used in proving our main result is that for a traffic network with multiple information types, if we reduce the total traffic demand, then there exists a route with strictly less flow. When some users have more information, they change their routing, redirecting it to some subnetwork A of the original network from some other subnetwork B (and since the original network is LI, both A and B are also LI). All else equal, this will increase flows in A and decrease flows in B. By the key property of LI networks, this will reduce flows in some route in B, and since users with more information have access to routes in B, this rerouting cannot increase their costs. Other users adjust their routing by allocating flows away from A (since flow in A has increased), which again by the LI property of the subnetwork implies that costs of some routes in A decrease, establishing the "if" part of our main result. The "only if" part is proved by showing that every non-SLI network embeds one of the collection of networks, and we demonstrate constructively that each one of these networks generates IBP (for some configuration of costs).

We should also note that, since SLI is a restrictive class of networks, and few real-world networks would fall into this class, we take this characterization to imply that the IBP is difficult to rule out in practice, and thus the new, highly anticipated route guidance technologies may make traffic problems worse.

Since the class of SLI networks plays a central role in our analysis, a natural question is whether identifying SLI networks is straightforward. We answer this question by showing that whether a given network is SLI or not can be determined in linear time. This result is based on the algorithms for identifying series-parallel networks proposed by Valdes et al. (1979), Schoenmakers (1995), and Eppstein (1992).

If, rather than considering a general change of information sets, we specialize the problem so that only one user type does not have complete information about the available set of routes and the change in question is to bring all users complete information, then we show that an IBP is possible if and only if the network is not *series parallel*. It is intuitive that this class of networks is less restrictive than SLI, since we are now considering a specific change in information sets (thus making IBP less likely to occur).

Our main focus is on traffic networks with a single origin–destination pair for which we provide a full characterization of network topologies for occurrence of IBP. In Section 6.4, we consider multiple origin–destination pairs and use our characterization to provide a sufficient condition on the network topology under which the IBP does not occur.

Our notion of the IBP closely relates to the classic Braess' paradox (BP), introduced in Braess (1968) and further studied in Murchland (1970) and Arnott and Small (1994), which considers whether an additional edge in the network can increase equilibrium cost. When BP occurs in a network, the IBP with a single information type also occurs (since the IBP with a single information type can be shown to be identical to BP). Various aspects of BP and congestion games in general are studied in Murchland (1970), Steinberg and Zangwill (1983), Dafermos and Nagurney (1984), Patriksson (1994), Bottom et al. (1999), Jahn et al. (2005), Ordóñez and Stier-Moses (2010), Meir and Parkes (2014), Nikolova and Stier-Moses (2015), Chen et al. (2015), and Feldman and Friedler (2015). Our characterization of the ICWE and IBP clarifies that our notion is different and, at least mathematically, more general. This can be seen readily from a comparison of our results to the most closely related papers to ours in the literature, Milchtaich (2005, 2006). The characterizations in Milchtaich (2006) imply that BP can be ruled out in series-parallel networks. Since the IBP is a generalization of BP, it should occur in a wider class of networks, and this is indeed what our result shows—SLI is a strict subset of series-parallel networks. This result also indicates that the IBP is a considerably more pervasive phenomenon than BP. Notably, the mathematical argument for our key theorem is different from Milchtaich (2006) because of the key difficulty relative to BP that not all users have access to the same set of edges, and thus changes in traffic that benefit some groups of users might naturally harm others by increasing the congestion on the routes that they were previously utilizing.

Issues related to Braess' paradox arise not only in the context of models of traffic but also in various models of communication, pricing and choice over congested goods, and electrical circuits. See, for example, Orda et al. (1993), Korilis et al. (1997), Kelly et al. (1998), and Low and Lapsley (1999) for communication networks; the classic works by Pigou (1920) and Samuelson (1952) as well as more recent works by Johari and Tsitsiklis (2004), Acemoglu and Ozdaglar (2007), Ashlagi et al. (2009), and Perakis (2004) for related economic problems; Frank (1981), Cohen and Horowitz (1991), and Cohen and Jeffries (1997) for mechanical systems and electrical circuits; and Rosenthal (1973) and Vetta (2002) for general game-theoretic approaches. This observation also implies that the results we present here are relevant beyond traffic networks, in fact to any resource allocation problem over a network subject to congestion considerations. As pointed out in Newell (1980) and Sheffi (1985), the Braess' paradox and related inefficiencies are a clear and present challenge to traffic engineers, who often try to restrict travel choices to improve congestion (e.g., via systems such as ramp metering on freeway entrances).

Other inefficiencies created by providing more information in the context of traffic networks have been studied in Mahmassani and Herman (1984), Ben-Akiva et al. (1991), Arnott et al. (1991), and Liu et al. (2016). In particular, Arnott et al. (1991) consider a model with atomic drivers in which users decide on their departure time and route choice. They show that providing imperfect information regarding capacity/delay of roads might be worse than providing no information. More broadly, inefficiencies created by providing more information in other contexts are studied in Maheswaran and Başar (2003), Sanghavi and Hajek (2004), Yang and Hajek (2005), Harel et al. (2014), Dughmi (2014), and Rogers et al. (2015), among others.

Because our analysis also presents "price of anarchy"–type results (i.e., bounds on the overall level of inefficiency that can occur in an ICWE), our paper is related to previous work on the price of anarchy in congestion and related games started by seminal works of Koutsoupias and Papadimitriou (1999) and Roughgarden and Tardos (2002) and followed by Correa et al. (2004), Correa et al. (2005), and Friedman (2004), as well as more generally to the analysis of equilibrium and inefficiency in the variants of this class of games, including Milchtaich (2004a, b), Acemoglu et al. (2007), Mavronicolas et al. (2007), Nisan et al. (2007), Arnott and Small (1994), Lin et al. (2004), Meir and Parkes (2015), and Anshelevich et al. (2008). Here, our result is that the presence of users with different information sets does not change the worst-case inefficiency traffic equilibrium as characterized, for example, in Roughgarden and Tardos (2002).

The rest of this paper is organized as follows. In Section 2, we introduce our model of traffic equilibrium with users that are heterogeneous in terms of the information about routes/edges they have access to, and then define the notion of information-constrained Wardrop equilibrium for this setting. In Section 3, we prove the existence and essential uniqueness of information-constrained Wardrop equilibrium. Before moving to our main focus, in Section 4, we review some graph-theoretic notions about series-parallel and linearly independent networks, and then we introduce the class of series of linearly independent networks and prove some basic properties of this class of networks, which are then used in the rest of our analysis. Section 5 defines our notion of the informational Braess' paradox. Section 6 contains our main result, showing that the informational Braess' paradox occurs "if and only if" the network is not in the class of series of linearly independent networks. Section 7 characterizes the worst-case inefficiency of information-constrained Wardrop equilibrium, and finally, Section 8 concludes. All the omitted proofs are included in the appendix.

2. Model

We first describe the environment and then introduce our notion of information-constrained Wardrop equilibrium.

2.1. Environment

We consider an undirected multigraph without self-loops denoted by $G = (V, \mathscr{E}, f)$ with vertex set V, edge set $\mathscr{E}$, and a function $f: \mathscr{E} \rightarrow \{\{u, v\}, u, v \in V, u \neq v\}$ that maps each edge to its end vertices. For ease of notation, we will refer to G as $(V, \mathscr{E})$ and denote an edge e with $f(e) = \{u, v\}$ by $e = (u, v)$. We use the terms "node" and "vertex" interchangeably. Each edge $e \in \mathscr{E}$ joins two (distinct) vertices u and v, referred to as the *end vertices* of e. An edge e and a vertex v are said to be *incident* to each other if v is an end vertex of e. A *path* $p \in G$ of length n ($n \geqslant 0$) is a sequence of edges $e_1, \ldots, e_n$ in $\mathscr{E}$, where e_i and e_{i+1} share a vertex. If an edge e appears on a path p, we write $e \in p$. The first and last vertices of a path p are called the initial and terminal vertices of p, respectively. If q is a path of the form $e_{n+1}, \ldots, e_m$, with the initial vertex the same as the terminal vertex of p but all the other vertices and edges of q do not belong to p, then $e_1, \ldots, e_n, e_{n+1}, \ldots, e_m$ is also a path, denoted by $p + q$. For a path p and two nodes v and u on it, we denote the *section* of path between u and v by p_{uv}.

Throughout the paper, we focus on an undirected multigraph $G = (V, \mathscr{E})$ together with an ordered pair of distinct vertices, called terminals, an origin O and a destination D, referred to as a *network*. A subnetwork of G is defined as $(V', \mathscr{E}')$, where $V' \subseteq V$ and $\mathscr{E}' \subseteq \mathscr{E}$ and for any $e = (u, v) \in \mathscr{E}'$, we have $u, v \in V'$. We assume that each vertex and edge belong to at least one path between the initial vertex O and the terminal vertex D. This assumption is without loss of generality because the vertices and edges that do not belong to any path from O to D are irrelevant for the purpose of sending traffic from O to D. Any path r with O as the initial vertex and D as the terminal vertex will be called a *route*. The set of all routes in a network is denoted by $\mathscr{R}$.

We suppose there are $K \geqslant 1$ types of users (we use the terms "users" and "players" interchangeably) and use the shorthand notation $[K] = \{1, \ldots, K\}$ to denote the set of types. Each type $i \in [K]$ has total *traffic demand* $s_i \in \mathbb{R}^+$, and we denote the vector of traffic demands by $s_{1:K} = (s_1, \ldots, s_K)$. For each type i, we use $\mathscr{E}_i \subseteq \mathscr{E}$ to denote the set of edges that type i knows and $\mathscr{R}_i$ to denote the routes formed by edges in $\mathscr{E}_i$ (assumed nonempty). We refer to $\mathscr{E}_i$ or $\mathscr{R}_i$ as type i's information set. We use $\mathscr{E}_{1:K} = (\mathscr{E}_1, \ldots, \mathscr{E}_K)$ to denote the information sets of all types.

We use $f^{(i)} = (f_r^{(i)}: r \in \mathscr{R}_i)$ to denote the flow vector of type i, where for all $r \in \mathscr{R}_i$, $f_r^{(i)} \geqslant 0$ represents the amount of traffic (flow) that type i sends on route r. We use $f^{(1:K)} = (f^{(1)}, \ldots, f^{(K)})$ to denote the flow vector of all types. Each edge of the network has a cost (latency) function $c_e: \mathbb{R}^+ \to \mathbb{R}^+$, which is continuous, nonnegative, and nondecreasing. We denote the set of all cost functions by $\mathbf{c} = \{c_e: e \in \mathscr{E}\}$. For instance, if all the cost functions are affine functions, then for any $e \in \mathscr{E}$, we would have $c_e(x) = a_e x + b_e$ for some $a_e, b_e \in \mathbb{R}^+$. We refer to $(G, \mathscr{E}_{1:K}, s_{1:K}, \mathbf{c})$ as a *traffic network with multiple information types*. A feasible flow is a flow vector $f^{(1:K)} = (f^{(1)}, \ldots, f^{(K)})$ such that for all $i \in [K]$, $f^{(i)}$ is a flow vector of type i (i.e., $f^{(i)}: \mathscr{R}_i \to \mathbb{R}^+$ and $\sum_{r \in \mathscr{R}_i} f_r^{(i)} = s_i$). We denote the total flow on each route r by f_r (i.e., $f_r = \sum_{i=1}^K f_r^{(i)}$).

2.2. Information-Constrained Wardrop Equilibrium

The cost of a route r with respect to a flow $(f^{(1)}, \ldots, f^{(K)})$ is the sum of the cost of the edges that belong to this route—that is, $c_r(f^{(1:K)}) = \sum_{e \in r} c_e(f_e)$, where f_e denotes the amount of traffic that passes through edge e (i.e., $f_e = \sum_{r \in \mathscr{R}: e \in r} f_r$).

We assume flows get allocated at equilibrium according to a "constrained" version of Wardrop's principle: flows of each user type are routed along routes in her information set with minimal (and hence equal) cost. We next formalize this equilibrium notion.

Definition 1 (ICWE). A feasible flow $f^{(1:K)} = (f^{(1)}, \ldots, f^{(K)})$ is an ICWE if for every $i \in [K]$ and every pair $r, \tilde{r} \in \mathscr{R}_i$ with $f_r^{(i)} > 0$, we have

$$c_r(f^{(1:K)}) \leqslant c_{\tilde{r}}(f^{(1:K)}). \quad (1)$$

Figure 1. Example of a Network with Edge Cost Functions Given by $c_{e_1}(x) = c_{e_4}(x) = c_{e_6}(x) = x$ and $c_{e_2}(x) = c_{e_5}(x) = c_{e_7}(x) = 1 + ax$ and $c_{e_3} = ax$ for Some $a > 0$

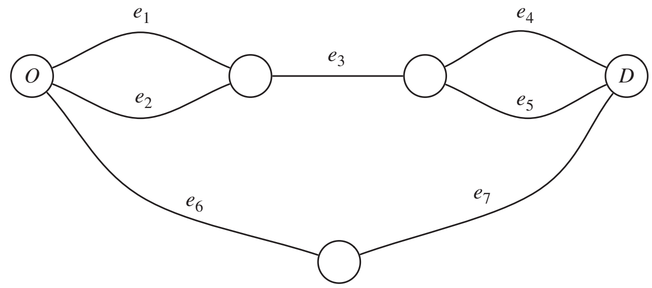

This implies that all the routes in $\mathscr{R}_i$ with positive flow from type i have the same cost, which is smaller or equal to the cost of any other route in $\mathscr{R}_i$. *The equilibrium cost of type i*, denoted by $c^{(i)}$, is then given by the cost of any route in $\mathscr{R}_i$ with positive flow from type i. Note that the Wardrop equilibrium (WE) is a special case of this definition for a traffic network with a single information type (i.e., $K = 1$).

We next provide an example that illustrates this definition and how it differs from the classic Wardrop equilibrium.

Example 1. Consider the network $G = (V, \mathscr{E})$ given in Figure 1 with $s_1 = s$, $s_2 = 1 - s$, and the cost functions specified in Figure 1. There are five different routes from origin to destination, which we denote by $r_1 = e_1e_3e_4$, $r_2 = e_1e_3e_5$, $r_3 = e_2e_3e_4$, $r_4 = e_2e_3e_5$, and $r_5 = e_6e_7$. We let $\mathscr{E}_1 = \mathscr{E}$ and $\mathscr{E}_2 = \{e_6, e_7\}$, which results in $\mathscr{R}_1 = \{r_1, r_2, r_3, r_4, r_5\}$ and $\mathscr{R}_2 = \{r_5\}$, respectively.

- If $s \leqslant (2+a)/(3+2a)$, the ICWE is $f_{r_1}^{(1)} = s$ and $f_{r_5}^{(2)} = 1 - s$. The equilibrium cost of type 1 is $c^{(1)} = c_{r_1}(f^{(1:2)}) = s + as + s = s(a+2)$. The equilibrium cost of type 2 is $c^{(2)} = c_{r_5}(f^{(1:2)}) = (1-s) + (1 + a(1-s)) = (1-s)(1+a) + 1$. Hence, the equilibrium cost of type 1 and type 2 users need not be the same.
- If $s > (2+a)/(3+2a)$, the ICWE is $f_{r_1}^{(1)} = (2+a)/(3+2a)$, $f_{r_5}^{(1)} = s - (2+a)/(3+2a) > 0$, and $f_{r_5}^{(2)} = 1 - s$, which gives $c^{(1)} = c^{(2)} = ((2+a)^2)/(3+2a)$. This illustrates that when different types use a common route in an equilibrium, their equilibrium costs are the same.

3. Existence of Information-Constrained Wardrop Equilibrium

In this section, we show that given a traffic network with multiple information types $(G, \mathscr{E}_{1:K}, s_{1:K}, \mathbf{c})$, an ICWE always exists and is "essentially" unique; that is, for each type, equilibrium cost is the same for all equilibria. Our proof of the existence and essential uniqueness of the ICWE relies on the following characterization, which is a straightforward extension of the well-known optimization characterization of the Wardrop equilibrium (see Beckmann et al. 1956, Smith 1979).

Proposition 1. *A flow $f^{(1:K)}$ is an ICWE if and only if it is a solution of the following optimization problem*:

$$\min \sum_{e\in\mathscr{E}} \int_0^{f_e} c_e(z)\,dz$$
$$f_e = \sum_{i=1}^{K} \sum_{r\in\mathscr{R}_i : e\in r} f_r^{(i)},$$
$$\sum_{r\in\mathscr{R}_i} f_r^{(i)} = s_i, \quad \text{and} \quad f_r^{(i)} \geqslant 0 \quad \text{for all } r\in\mathscr{R}_i. \tag{2}$$

We call $\sum_{e\in\mathscr{E}} \int_0^{f_e} c_e(z)\,dz$ the potential function and denote it by Φ.

Using the characterization of ICWE as the minimizer of a potential function, we can now show the existence and essential uniqueness.

Theorem 1 (Existence and Uniqueness of ICWE). *Let $(G, \mathscr{E}_{1:K}, s_{1:K}, \mathbf{c})$ be a traffic network with multiple information types.*

- *There exists an ICWE flow $f^{(1:K)} = (f^{(1)}, \ldots, f^{(K)})$.*
- *The ICWE is essentially unique in the sense that if $f^{(1:K)}$ and $\tilde{f}^{(1:K)}$ are both ICWE flows, then $c_e(f_e) = c_e(\tilde{f}_e)$ for every edge $e\in\mathscr{E}$.*

Remark 1. As shown in Milchtaich (2005), Gairing et al. (2006), and Mavronicolas et al. (2007) the essential uniqueness of equilibrium does not hold for multiple-type congestion games where different types of users have different cost functions for the same edge. This class of congestion games is also referred to as *player-specific congestion games*. Several conditions on the edge cost functions and network topology have been proposed to guarantee the existence of an essentially unique equilibrium (see Konishi et al. 1997, Voorneveld et al. 1999, Milchtaich 2005, Mavronicolas et al. 2007, Georgiou et al. 2009, Gairing and Klimm 2013). In particular, Milchtaich (2005) provides sufficient and necessary conditions on the network topology under which an essentially unique equilibrium exists. Mavronicolas et al. (2007) and Georgiou et al. (2009) show that when the edge costs are affine functions and differ by a player-specific additive constant, then an equilibrium exists. Our model is a special case of a player-specific congestion game in which the cost of an edge e for a type i user is $c_e(\cdot)$ if $e\in\mathscr{E}_i$ and ∞ otherwise. Therefore, the results of Mavronicolas et al. (2007) and Georgiou et al. (2009) can directly be used to establish the existence of an equilibrium in our model. For completeness, we provide an alternative proof of Theorem 1 in Appendix A.1 based on the classical results of Beckmann et al. (1956), Schmeidler (1973), Smith (1979), and Milchtaich (2000).

Theorem 1 assumes that the cost functions are nondecreasing. If we strengthen this assumption to strictly increasing cost functions, then the results of Roughgarden and Tardos (2002), Mavronicolas et al. (2007), and Georgiou et al. (2009) show that the essential uniqueness result can be strengthened. In this case, the total flow on any edge at any equilibrium would be the same.

4. Some Graph-Theoretic Notions

In this section, we first present two classes of networks—namely, series-parallel and linearly independent networks—which we use in our characterization of an IBP. In preparation for our main graph-theoretic results, we also present equivalent characterizations of these networks and delineate the relations among them. Finally, we define a new class of networks termed series of linearly independent and present a characterization for it in terms of embedding of a few basic networks.

Definition 2 (Series-Parallel Network (SP)). A (two-terminal) network is called series parallel if two routes never pass through an edge in opposite directions. Equivalently, as was shown by Riordan and Shannon (1942), a network is series parallel if and only if

(i) it comprises a single edge between O and D, or

(ii) it is constructed by connecting two series-parallel networks in series (i.e., by joining the destination of one series-parallel network with the origin of the other one), or

(iii) it is constructed by connecting two series-parallel networks in parallel (i.e., by joining the origins and destinations of two series-parallel networks).

As an example, the networks shown in Figures 2(b) and 2(c) are series-parallel networks, while the network shown in Figure 2(a) is not. The reason is that two routes $e_1e_5e_4$ and $e_2e_5e_3$ pass through the edge e_5 in opposite directions.

An important subclass of series-parallel networks are *linearly independent* networks.

Definition 3 (Linearly Independent Network). A (two-terminal) network is called linearly independent if each route has at least one edge that does not belong to any other route. Equivalently, as was shown by Holzman and Law-yone (2003), a network is linearly independent if and only if

(i) it comprises a single edge between O and D, or

(ii) it is constructed by connecting a linearly independent network in series with a single edge network, or

(iii) it is constructed by connecting two linearly independent networks in parallel.

This class is termed linearly independent because of an algebraic characterization of the routes when viewed as vectors in the edge space. In particular, for any $r\in\mathscr{R}$, let $\mathbf{v}_r \in \mathbb{F}_2^{|\mathscr{E}|}$ be $\mathbf{v}_r = (v_r^1, \ldots, v_r^{|\mathscr{E}|})$, where $v_r^i = 1$ if $e_i \in r$ and 0 otherwise. A network G is LI if and only

Figure 2. Networks That Cannot Be Embedded in SP and LI Networks: Network (a) Is Not Embedded in SP Networks; Networks (a), (b), and (c) Are Not Embedded in LI Networks

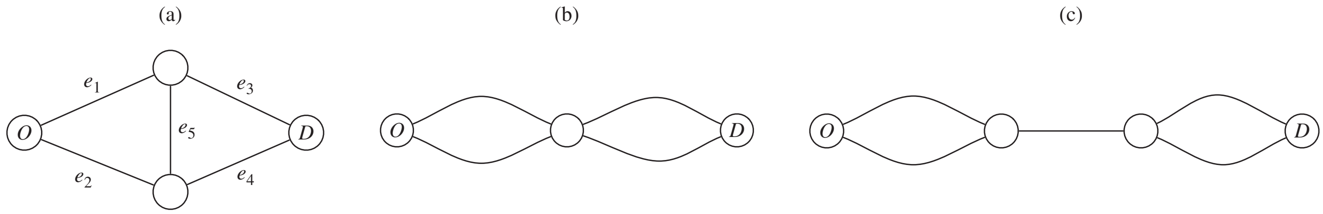

if the set of vectors $\{\mathbf{v}_r: r \in \mathcal{R}\}$ is linearly independent (see Milchtaich 2006, Diestel 2000).

As Definitions 2 and 3 make clear, the class of linearly independent networks is a subset of the class of series-parallel networks. An alternative characterization of linearly independent and series-parallel networks is based on the "graph embedding" notion, shown by Duffin (1965) and Milchtaich (2006), respectively. We next define a graph embedding and then present these characterizations, which will be used later in our analysis.

Definition 4 (Embedding). A network H is embedded in the network G if we can start from H and construct G by applying the following steps in any order:

(i) Divide an edge; that is, replace an edge with two edges with a single common end node.
(ii) Add an edge between two nodes.
(iii) Extend origin or destination by one edge.

Proposition 2. (a) (*Milchtaich* 2006): *A network G is LI if and only if none of the networks shown in Figure 2 are embedded in it. Furthermore, a network G is LI if and only if for every pair of routes r and r' and every vertex $v \neq O, D$ common to both routes, either the section r_{Ov} is equal to r'_{Ov} or the section r_{vD} is equal to r'_{vD}.*

(b) (*Duffin* 1965 *and Milchtaich* 2006): *A network G is SP if and only if the network shown in Figure 2(a) is not embedded in it. Furthermore, a network G is SP if and only if the vertices can be indexed in such a way that, along each route, they have increasing indices.*

This proposition shows that series-parallel networks are those in which the network shown in Figure 2(a), which is referred to as the *Wheatstone network* (see Braess 1968), is not embedded. LI networks, in addition, also exclude embeddings of series-parallel networks that have routes that "cross" as indicated in Figures 2(b) and 2(c).

We now introduce a new class of networks, which we refer to as *series of linearly independent networks*.

Definition 5 (Series of Linearly Independent Network). A (two-terminal) network G is called series of linearly independent if and only if

(i) it comprises a single linearly independent network, or

(ii) it is constructed by connecting two SLI networks in series.

A biconnected LI network is called an *LI block*, where a graph is biconnected if it is connected and after removing any node and its incident edges the graph remains connected (see Bondy and Murty 1976, chap. 3). Equivalently, a network G is SLI if and only if it is constructed by connecting several LI blocks in series (see Appendix A.2.1 for a formal proof). We refer to each of these blocks as an LI block of SLI network G.

We next provide a new characterization of SLI networks in terms of graph embedding using the characterizations for SP and LI networks presented in Proposition 2.

Theorem 2 (Characterization of SLI). *A network G is SLI if and only if none of the networks shown in Figure 3 are embedded in it.*

The class of SLI networks is a subset of series-parallel networks and a superset of linearly independent networks. This class plays an important role in our characterization of networks that exhibit IBP. Valdes et al. (1979) provided an algorithm to determine whether a given network is SP in $O(|\mathscr{E}| + |V|)$ steps based on a tree decomposition of SP networks. This leads to the question whether one can find a linear time algorithm (i.e., linear in the number of vertices and edges) to recognize an SLI network. We next use the results of Valdes et al. (1979) to show that we can recognize whether a given network is SLI in linear time.

Proposition 3. *There exists an algorithm that can determine whether a given network G is SLI in $O(|\mathscr{E}| + |V|)$.*

5. Informational Braess' Paradox

We first present the classical BP, which is defined for a traffic network with a single type of user with $\mathscr{E}_1 = \mathscr{E}$, denoted by $(G, \mathscr{E}_1, s_1, \mathbf{c})$.

Definition 6 (BP). Consider a traffic network with a single information type $(G, \mathscr{E}_1, s_1, \mathbf{c})$. BP occurs if there exists another set of cost functions $\hat{\mathbf{c}}$ with $\hat{c}_e(x) \leqslant c_e(x)$ for all $e \in \mathscr{E}$ and $x \in \mathbb{R}^+$, such that the equilibrium cost of $(G, \mathscr{E}_1, s_1, \hat{\mathbf{c}})$ is strictly larger than the equilibrium cost of $(G, \mathscr{E}_1, s_1, \mathbf{c})$.

Figure 3. Networks That Cannot Be Embedded in SLI Networks

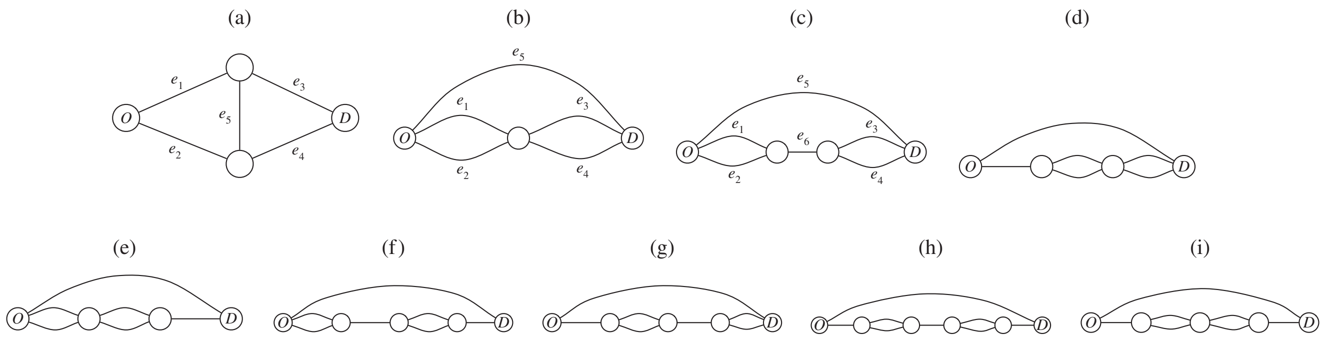

BP refers to an unexpected increase in equilibrium cost in response to a decrease in edge costs. We next discuss the IBP, which arises when providing more information to a subset of users in a traffic network increases those users' costs.

Definition 7 (IBP). Consider a traffic network with multiple information types $(G, \mathscr{E}_{1:K}, s_{1:K}, \mathbf{c})$. IBP occurs if there exist expanded information sets $\tilde{\mathscr{E}}_{1:K}$ with $\mathscr{E}_1 \subset \tilde{\mathscr{E}}_1 \subseteq \mathscr{E}$ and $\tilde{\mathscr{E}}_i = \mathscr{E}_i$, for $i = 2, \ldots, K$, such that the equilibrium cost of type 1 in $(G, \tilde{\mathscr{E}}_{1:K}, s_{1:K}, \mathbf{c})$ is strictly larger than the equilibrium cost of type 1 in $(G, \mathscr{E}_{1:K}, s_{1:K}, \mathbf{c})$. We denote the equilibrium cost of type $i \in [K]$ before and after the expansion of information sets by $c^{(i)}$ and $\tilde{c}^{(i)}$, respectively.

The choice of type 1 users in this definition is without loss of generality; that is, we assume that the information set of only one type expands and the information sets of the rest of the types remain the same. In comparing IBP to BP, first note that BP occurs in a network if and only if a special case of BP occurs in which we decrease the cost of one of the edges from infinity to its actual cost; that is, equilibrium cost increases by adding a new edge to the network. The "if" part holds by definition, and the "only if" part holds because the special case of BP occurs in the Wheatstone network (as presented in Example 2(a)) and the Wheatstone network is embedded in any network that features BP as shown by Milchtaich (2005). In light of this, it follows that the occurrence of IBP is a generalization of that of BP since addition of a new edge to the network can be viewed as expansion of the information set of a type to include that edge in a traffic network with a single information type.

The next example shows that IBP occurs in all networks shown in Figure 3—that is, all the basic networks that are embedded in non-SLI networks.

Example 2. In this example we will show that for all networks shown in Figure 3, there exists an assignment of cost functions along with information sets for which IBP occurs.

(a) IBP occurs for the Wheatstone network shown in Figure 3(a). This follows from the occurrence of BP on the Wheatstone network as shown in Braess (1968). We will provide the example for the sake of completeness in Appendix A.3.1.

(b) Consider the network shown in Figure 3(b) with cost functions given by $c_{e_1}(x) = \frac{1}{2}x$, $c_{e_2}(x) = x + \frac{3}{4}$, $c_{e_3}(x) = \frac{4}{3}x$, $c_{e_4}(x) = 2$, and $c_{e_5}(x) = x$. The information sets are $\mathscr{E}_1 = \{e_2, e_3, e_5\}$, $\mathscr{E}_2 = \{e_1, e_4, e_5\}$, and $\tilde{\mathscr{E}}_1 = \{e_1, e_2, e_3, e_5\}$. For $s_1 = 13/4$ and $s_2 = 1$, the equilibrium flows are

$$f^{(2)}_{e_1e_4} = 1, \quad f^{(2)}_{e_5} = 0, \quad f^{(1)}_{e_2e_3} = \tfrac{3}{4}, \quad f^{(1)}_{e_5} = \tfrac{10}{4},$$
$$\tilde{f}^{(2)}_{e_1e_4} = 0, \quad \tilde{f}^{(2)}_{e_5} = 1, \quad \tilde{f}^{(1)}_{e_2e_3} = 0, \quad \tilde{f}^{(1)}_{e_1e_3} = \tfrac{6}{4}, \quad \tilde{f}^{(1)}_{e_5} = \tfrac{7}{4}.$$

The resulting equilibrium costs are $c^{(1)} = c^{(2)} = 10/4$ and $\tilde{c}^{(1)} = \tilde{c}^{(2)} = 11/4$. Since $\tilde{c}^{(1)} > c^{(1)}$, IBP occurs in this network. The main intuition for this example is as follows. After adding e_1 to $\mathscr{E}_1$, type 1 users will no longer use e_2e_3 and instead redirect part of their flow over e_1e_3. This, in turn, will increase the cost of e_1e_4 for type 2 users and induce them to redirect all their flow from e_1e_4 to e_5. In balancing the costs of e_1e_3 and e_5 for type 1 users, their equilibrium cost goes up.

(c) Finally, for the networks shown in Figures 3(c)–3(i), IBP occurs if we use the same setting as in part (b) and include extra edges in all information sets with zero cost.

Remark 2. In Appendix A.3.2, we show that Example 2(b) is not degenerate and provide an infinite set of (affine) cost functions for which IBP occurs in this network. Similar to Example 2(c), this argument extends to show that there are infinitely many cost functions for which IBP occurs in networks shown in Figures 3(c)–3(i). Finally, for the network shown in Figure 3(a), there are infinitely many cost functions for which BP occurs when edge e_5 is added; hence IBP occurs as well (see, e.g., Steinberg and Zangwill 1983).

In a seminal paper, Milchtaich (2006) provided necessary and sufficient conditions on the network topology under which BP occurs. In particular, Milchtaich (2006) showed that for a given traffic network with a single information type $(G, \mathscr{E}_1, s_1, \mathbf{c})$, BP does not occur if and only if G is SP. That is, if G is SP, then for any assignment of cost functions $\mathbf{c}$ and traffic demand, BP does not occur, and if G is not SP, then there exists an assignment of cost functions $\mathbf{c}$ for which BP occurs.

We next investigate conditions on the topology of the network under which IBP occurs. In a way similar to the characterization provided by Milchtaich (2006), we will identify classes of networks for which IBP does not occur regardless of the cost functions of the edges. Since, as already noted, IBP is a strict generalization of BP, we will see that IBP can occur in a broader class of networks, underscoring the problem mentioned in the introduction that IBP is likely to be a more pervasive problem.

6. Characterization of Informational Braess' Paradox

In this section, after establishing the key lemmas that underpin the rest of our analysis, we provide our main characterization of IBP. In Section 6.3 we provide a characterization for IBP for a more restricted type of change in information sets. We conclude this section with a discussion of extensions of our results to multiple origin–destination pairs.

6.1. Three Key Lemmas

The following lemmas identify properties of the traffic network consisting of heterogeneous users over an LI network.

Lemma 1. (a) *Given an LI network G, let $f^{(1:K)}$ and $\tilde{f}^{(1:K)}$ be two arbitrary nonidentical feasible flows for two traffic networks $(G, \mathscr{E}_{1:K}, s_{1:K}, \mathbf{c})$ and $(G, \mathscr{E}_{1:K}, \tilde{s}_{1:K}, \mathbf{c})$, respectively. If $\sum_{i=1}^{K} s_i \geqslant \sum_{i=1}^{K} \tilde{s}_i$, then there exists a route r such that $\sum_{i=1}^{K} f_r^{(i)} > \sum_{i=1}^{K} \tilde{f}_r^{(i)}$ and $f_e \geqslant \tilde{f}_e$ for all $e \in r$.*

(b) *Given an LI network G, let $c^{(i)}$ and $\tilde{c}^{(i)}$ denote the equilibrium cost of type $i \in [K]$ users in traffic networks $(G, \mathscr{E}_{1:K}, s_{1:K}, \mathbf{c})$ and $(G, \tilde{\mathscr{E}}_{1:K}, s_{1:K}, \mathbf{c})$, respectively. If $\mathscr{E}_1 \subseteq \tilde{\mathscr{E}}_1$ and $\tilde{\mathscr{E}}_i = \mathscr{E}_i$, for $i = 2, \ldots, K$, then there exists some $i \in [K]$ such that $\tilde{c}^{(i)} \leqslant c^{(i)}$.*

This lemma directly follows from Milchtaich (2006, lemma 5 and theorem 3). The first part of the lemma shows that in an LI network, if the total traffic increases, then there exists at least one route whose flow strictly increases, and the flow on each of its edges weakly increases. The second part shows that in an LI network, if we expand the information set of type 1 users, then the equilibrium cost of at least one of the types does not increase. In fact, a similar argument shows that even if we expand the information set of multiple types, then the equilibrium cost of at least one of the types does not increase (see Milchtaich 2006, theorem 3). Note that this result is not sufficient for establishing that IBP does not occur over LI networks because what we need to establish is that it is the equilibrium cost of type 1 users that does not increase. For completeness, in Appendix A.4.1, we show how this lemma follows from the results of Milchtaich (2006).

The next lemma shows a property of equilibrium flows and equilibrium costs in a network, which is the result of attaching two networks in series. We use the following definition to state the lemma. Suppose $f^{(1:K)}$ is a feasible flow for $(G, \mathscr{E}_{1:K}, \mathbf{s}_{1:K}, \mathbf{c})$, where G is the result of attaching G_1 and G_2 in series. We denote the attaching point of G_1 and G_2 by D_1. The *restriction* of $f^{(1:K)}$ to G_1 (similarly to G_2) is defined as $\bar{f}^{(1:K)} = (\bar{f}^{(1)}, \ldots, \bar{f}^{(K)})$, where the flow of type i on any route $\bar{r}$ in G_1 is the summation of the flows of type i on all routes of G that contain $\bar{r}$. Formally, for any $i \in [K]$, we have $\bar{f}^{(i)}(\bar{r}) = \sum_{r \in \bar{\mathscr{R}}_i(\bar{r})} f^{(i)}(r)$, where $\bar{\mathscr{R}}_i(\bar{r}) = \{r \in \mathscr{R}_i\colon r_{OD_1} = \bar{r}\}$. Note that $\bar{f}^{(1:K)}$ is a feasible flow on G_1.

Lemma 2. (a) *If G is the result of attaching two networks G_1 and G_2 in series, then the restriction of an equilibrium flow for G to each of G_1 and G_2 is an equilibrium flow.*

(b) *If G is the result of attaching two networks G_1 and G_2 in series, then the equilibrium cost of any type on G is the summation of the equilibrium costs of that type on G_1 and G_2.*

The third lemma shows our key lemma that we will use in the proof of Theorem 3. Intuitively, this lemma states that in an LI network, if we decrease the traffic on one subset of routes $\mathscr{R}_A$ of the network and reroute it through the rest of the routes in the network, denoted by $\mathscr{R}_B = \mathscr{R} \backslash \mathscr{R}_A$, then the maximum cost improvement over all the routes in $\mathscr{R}_A$ cannot be smaller than the minimum cost improvement over all the routes in $\mathscr{R}_B$. This result will enable us to establish that in an LI or SLI network, the reallocation of traffic because of one type of user obtaining more information cannot harm that type.

Lemma 3. *Given an LI network G, we let $\mathscr{R}_A, \mathscr{R}_B \neq \varnothing$ denote a partition of routes $\mathscr{R}$ (i.e., $\mathscr{R}_B = \mathscr{R} \backslash \mathscr{R}_A$). We let $f^{(1:K)}$ and $\tilde{f}^{(1:K)}$ be two feasible flows for traffic networks $(G, \mathscr{E}_{1:K}, s_{1:K}, \mathbf{c})$ and $(G, \mathscr{E}_{1:K}, \tilde{s}_{1:K}, \mathbf{c})$, respectively. For these two flows, we let the traffic over $\mathscr{R}_A$ and $\mathscr{R}_B$ be $s_A = \sum_{r \in \mathscr{R}_A} \sum_{i=1}^{K} f_r^{(i)}$, $\tilde{s}_A = \sum_{r \in \mathscr{R}_A} \sum_{i=1}^{K} \tilde{f}_r^{(i)}$, $s_B = \sum_{r \in \mathscr{R}_B} \sum_{i=1}^{K} f_r^{(i)}$, and $\tilde{s}_B = \sum_{r \in \mathscr{R}_B} \sum_{i=1}^{K} \tilde{f}_r^{(i)}$. If $\tilde{s}_A \leqslant s_A$ and $\tilde{s}_B \geqslant s_B$, then we have*

$$\max_{r \in \mathscr{R}_A} \{c_r - \tilde{c}_r\} \geqslant \min_{r \in \mathscr{R}_B} \{c_r - \tilde{c}_r\},$$

where for any route r, c_r and $\tilde{c}_r$ denote the cost of this route with flows $f^{(1:K)}$ and $\tilde{f}^{(1:K)}$, respectively.

Before proving this lemma for a general LI network, let us show it for the special case where G consists of parallel edges from O to D. In this case $\mathscr{R}_A$ and $\mathscr{R}_B$ are two disjoint sets of edges from O to D. Since $\tilde{s}_A \leqslant s_A$, there exists an edge e_A in $\mathscr{R}_A$ such that $\tilde{f}_{e_A} \leqslant f_{e_A}$.

Similarly, since $\tilde{s}_B \geqslant s_B$, there exists $e_B \in \mathcal{R}_B$ such that $\tilde{f}_{e_B} \geqslant f_{e_B}$. Since the cost functions are nondecreasing, we have

$$\max_{r\in\mathcal{R}_A}\{c_r-\tilde{c}_r\} \geqslant c_{e_A}(f_{e_A})-c_{e_A}(\tilde{f}_{e_A}) \geqslant 0 \geqslant c_{e_B}(f_{e_B})-c_{e_B}(\tilde{f}_{e_B}) \geqslant \min_{r\in\mathcal{R}_B}\{c_r-\tilde{c}_r\},$$

which is the desired result. The proof for the general case is by induction on the number of edges and is included next.

Proof. We first note a consequence of Proposition 2.

Claim 1. *If a network G is LI then for any vertex v, either the sections from O to v of all routes that pass through v (which consists of v and all the vertices and edges preceding it on the route) are identical or the sections from v to D of all routes that pass through v (which consists of v and all the vertices and edges succeeding it on the route) are identical. Consider a route r that passes through v. First, note that since G is SP, part* (b) *of Proposition* 2 *implies that the only common node of* r_{Ov} *and* r_{vD} *is v. Also, the only common node of* r'_{Ov} *and* r'_{vD} *is v. Claim* 1 *follows since if the contrary holds, then there exist two routes* $r = r_{Ov} + r_{vD}$ *and* $r' = r'_{Ov} + r'_{vD}$ *with a common vertex v such that* $r_{Ov} \neq r'_{Ov}$ *and* $r_{vD} \neq r'_{vD}$*. This contradicts the statement of part* (a) *of Proposition* 2.

We now prove Lemma 3 using induction on the number of edges. For a single edge, it evidently holds. For a general LI network, we have the following cases:

- There exists $r \in \mathcal{R}_A$ such that $c_r \geqslant \tilde{c}_r$ and $r' \in \mathcal{R}_B$ such that $c_{r'} \leqslant \tilde{c}_{r'}$. This leads to

$$\max_{r\in\mathcal{R}_A}\{c_r-\tilde{c}_r\} \geqslant c_r-\tilde{c}_r \geqslant 0 \geqslant c_{r'}-\tilde{c}_{r'} \geqslant \min_{r\in\mathcal{R}_B}\{c_r-\tilde{c}_r\},$$

which concludes the proof in this case.

- For any $r \in \mathcal{R}_A$, we have $c_r < \tilde{c}_r$. We break the proof into three steps.

Step 1. There exist a route $r \in \mathcal{R}_A$ and an edge $e \in r$ with the following properties: (i) The flow on r from $\tilde{s}_A$ is less than or equal to the flow on r from s_A. (ii) The flow on e from $\tilde{s}_B$ is larger than the flow on e from s_B, and the flow on e from $\tilde{s}_A$ is less than or equal to the flow on e from s_A.

This step follows from invoking part (a) of Lemma 1. Since $\tilde{s}_A \leqslant s_A$, using part (a) of Lemma 1, there exists a route $r \in \mathcal{R}_A$ such that the flow on each edge of r from $\tilde{s}_A$ is less than or equal to the flow from s_A. However, we know that the overall cost of any $r \in \mathcal{R}_A$ has gone up (i.e., $\tilde{c}_r > c_r$). This implies that there exists an edge $e \in r$ such that the flow from $\tilde{s}_B$ on e is more than the flow from s_B on e.

Step 2. Let $\mathcal{R}_e$ denote the set of routes using edge $e = (u_e, v_e)$ as defined in Step 1. Note that $\mathcal{R}_e$ has the following properties: (i) Either there exists a vertex $D' \in V$ such that all routes $r \in \mathcal{R}_e$ have a common path from O to v_e and a common path from D' to D or there exists a vertex $O' \in V$ such that all routes $r \in \mathcal{R}_e$ have a common path from u_e to D and a common path from O to O'. Without loss of generality, we assume it is the former case. (ii) There exists a subnetwork G' with origin $O' = v_e$ and destination D' such that for the restricted parts of $\mathcal{R}_A$ and $\mathcal{R}_B$ over G', denoted by $\mathcal{R}'_A$ and $\mathcal{R}'_B$, if we let s'_A, $\tilde{s}'_A$, s'_B, and $\tilde{s}'_B$ to denote the corresponding traffic demands on $\mathcal{R}'_A$ and $\mathcal{R}'_B$, then we have $\tilde{s}'_A \leqslant s'_A$ and $\tilde{s}'_B \geqslant s'_B$.

Using Claim 1, for an edge $e = (u_e, v_e)$, either there is a unique path from O to v_e or there is a unique path from v_e to D; we assume without loss of generality it is the former case. We let D' be the first node on route r such that all routes in $\mathcal{R}_e$ coincide from D' to D. Therefore, all routes $r \in \mathcal{R}_e$ have a common path from O to v_e and a common path from D' to D, showing the first property.

We next show that the subnetwork consisting of all routes from v_e to D', denoted by G', satisfies the second property. To see this, note that the flows on G' are only the ones that are passing through edge e. From Step 1, we know that the flow on e from $\tilde{s}_B$ is larger than the flow on e from s_B and the flow on e from $\tilde{s}_A$ is less than or equal to the flow on e from s_A, showing the second property.

Step 3. Using Steps 1 and 2 and the induction hypothesis for G', we will show that

$$\max_{r\in\mathcal{R}_A}\{c_r-\tilde{c}_r\} \geqslant \min_{r\in\mathcal{R}_B}\{c_r-\tilde{c}_r\}.$$

First note that $\mathcal{R}_e \cap \mathcal{R}_A \neq \emptyset$, since $r \in \mathcal{R}_A$ and $e \in r$. Furthermore, $\mathcal{R}_e \cap \mathcal{R}_B \neq \emptyset$, since, as explained in Step 1, the flow on e from $\tilde{s}_B$ is strictly positive. This, in turn, shows that $\mathcal{R}'_A$ and $\mathcal{R}'_B$ are nonempty. Using Step 2, all the conditions of Lemma 3 hold for subnetwork G'. Therefore, we can use the induction hypothesis for LI network G' to obtain

$$\max_{r\in\mathcal{R}'_A}\{c_r-\tilde{c}_r\} \geqslant \min_{r\in\mathcal{R}'_B}\{c_r-\tilde{c}_r\}.$$

Using Step 2, for all the routes in $\mathcal{R}_e$ the costs of going from O to O', denoted by $c_{O'\to O}$, are the same. Similarly, the costs of all routes in $\mathcal{R}_e$ going from D' to D, denoted by $c_{D'\to D}$, are the same.

Therefore, we have

$$\begin{aligned}\max_{r\in\mathcal{R}_A}\{c_r-\tilde{c}_r\} &\geqslant \max_{r\in\mathcal{R}_A\cap\mathcal{R}_e}\{c_r-\tilde{c}_r\} = (c_{O\to O'}-\tilde{c}_{O\to O'})\\ &\quad+\max_{r\in\mathcal{R}'_A}\{c_r-\tilde{c}_r\}+(c_{D'\to D}-\tilde{c}_{D'\to D})\\ &\geqslant (c_{O\to O'}-\tilde{c}_{O\to O'})+\min_{r\in\mathcal{R}'_B}\{c_r-\tilde{c}_r\}+(c_{D'\to D}-\tilde{c}_{D'\to D})\\ &= \min_{r\in\mathcal{R}_B\cap\mathcal{R}_e}\{c_r-\tilde{c}_r\} \geqslant \min_{r\in\mathcal{R}_B}\{c_r-\tilde{c}_r\},\end{aligned}$$

which concludes the proof in this case.

- For any $r \in \mathcal{R}_B$, we have $c_r > \tilde{c}_r$. The proof of this case is similar to the previous case. We state the three steps without repeating the reasoning for each of them.

Step 1. There exists a route $r \in \mathcal{R}_B$ and an edge $e \in r$ with the following properties: (i) The flow on r from $\tilde{s}_B$ is larger than or equal to the flow on r from s_B. (ii) The flow on e from $\tilde{s}_A$ is smaller than the flow on e from s_A, and the flow on e from $\tilde{s}_B$ is larger than or equal to the flow on e from s_B.

Step 2. Let $\mathcal{R}_e$ denote the set of routes using edge $e = (u_e, v_e)$ as defined in Step 1. We have the following properties: (i) Either there exists a vertex $D' \in V$ such that all routes $r \in \mathcal{R}_e$ have a common path from O to v_e and a common path from D' to D or there exists a vertex $O' \in V$ such that all routes $r \in \mathcal{R}_e$ have a common path from u_e to D and a common path from O to O'. Without loss of generality, we assume it is the former case. (ii) There exists a subnetwork G' with origin $O' = v_e$ and destination D' such that for the restricted parts of $\mathcal{R}_A$ and $\mathcal{R}_B$ over G', denoted by $\mathcal{R}'_A$ and $\mathcal{R}'_B$, if we let s'_A, $\tilde{s}'_A$, s'_B, and $\tilde{s}'_B$ denote the corresponding traffic on $\mathcal{R}'_A$ and $\mathcal{R}'_B$, then we have $\tilde{s}'_A \leqslant s'_A$ and $\tilde{s}'_B \geqslant s'_B$.

Step 3. Again, using Steps 1 and 2 and the induction hypothesis for G', we have

$$\max_{r \in \mathcal{R}_A} \{c_r - \tilde{c}_r\} \geqslant \min_{r \in \mathcal{R}_B} \{c_r - \tilde{c}_r\},$$

which completes the proof.

6.2. Characterization of the Informational Braess' Paradox

We next present our main result, which states that IBP does not occur if and only if the network is SLI. The idea of this result, as already discussed in the introduction, is the following. To show the "if" part, we note that using Lemma 2 it suffices to show IBP does not occur in LI networks. Consider an expansion of the information set of type 1 and the new equilibrium flows. If the equilibrium cost of type 1 increases, then ICWE definition implies the following:

• Consider all types with increased equilibrium costs (including type 1). All routes used by these types (in the equilibrium before information expansion) have higher costs in the new equilibrium.

• Consider all types with decreased equilibrium costs. All routes used by these types (in the equilibrium after information expansion) have lower costs in the new equilibrium.

Using these two claims, it follows that the total flow sent over the routes with higher costs is lower, and the total flow sent over the routes with lower costs is higher (see Figure 4). Since the network is LI, Lemma 3 leads to a contradiction. The "only if" part holds because any non-SLI network embeds one of the networks shown in Figure 3, and an IBP can be constructed for each of them (Example 2), which then extends to an IBP for the non-SLI network.

Theorem 3 (Characterization of IBP). *IBP does not occur if and only if G is SLI. More specifically, we have the following*:

(a) *If G is SLI, for any traffic network $(G, \mathcal{E}_{1:K}, s_{1:K}, \mathbf{c})$ with arbitrary assignment of cost functions $\mathbf{c}$, K, traffic demands $s_{1:K}$, and information sets $\mathcal{E}_{1:K}$, IBP does not occur.*

(b) *If G is not SLI, there exists an assignment of cost functions $\mathbf{c}$, K, traffic demands $s_{1:K}$, and information sets $\mathcal{E}_{1:K}$ in which IBP occurs.*

Proof of Part (a). To reach a contradiction, suppose that $\tilde{c}^{(1)} > c^{(1)}$. By Definition 5, G is obtained from attaching several LI blocks in series, denoted by $G_1, \ldots, G_N$ for some $N \geqslant 1$. Using part (b) of Lemma 2,

Figure 4. (Color online) Proof of Theorem 3: Set A (B) Represents Types with Higher (Lower) Equilibrium Costs and Set $\mathcal{R}_A$ ($\mathcal{R}_B$) Represents Routes with Higher (Lower) Costs

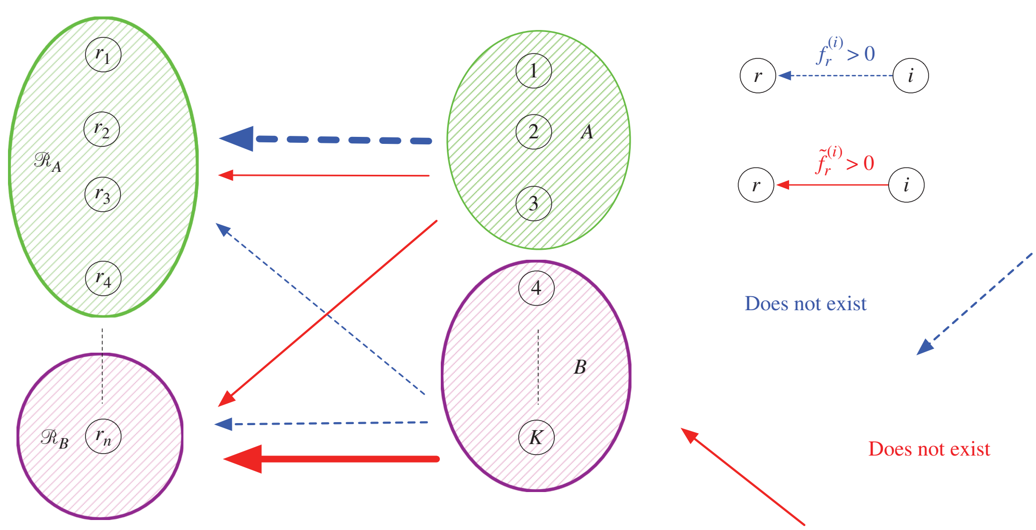

Note. There is no dashed (blue) arrow from A to $\mathcal{R}_B$, which illustrates Claim 1, and there is no solid arrow (red) from B to $\mathcal{R}_A$, which illustrates Claim 2.

we have $\tilde{c}^{(1)} = \sum_{t=1}^{N} \tilde{c}_t^{(1)} > \sum_{t=1}^{N} c_t^{(1)} = c^{(1)}$, where $c_t^{(1)}$ denotes the equilibrium cost of type 1 users in G_t. Therefore, there exists one LI block such as j for which $\tilde{c}_j^{(1)} > c_j^{(1)}$. Also, using part (a) of Lemma 2, the restriction of equilibrium flows $f^{(1:K)}$ and $\tilde{f}^{(1:K)}$ to G_j creates an equilibrium flow for this LI block. Therefore, IBP occurs in LI block G_j. In the rest of the proof of part (a), we will assume IBP occurs in an LI block (and hence LI network) and reach a contradiction. We let $f^{(1:K)}$ and $\tilde{f}^{(1:K)}$ be the equilibrium flows before and after the information set expansion, respectively. Also, for any route $r \in \mathcal{R}$, we let c_r and $\tilde{c}_r$ denote the cost of route r with flows $f^{(1:K)}$ and $\tilde{f}^{(1:K)}$, respectively.

We partition the set $[K]$ into groups A and B as follows:

$$A = \{i \in [k]: \tilde{c}^{(i)} > c^{(i)}\},$$

and

$$B = \{i \in [k]: \tilde{c}^{(i)} \leqslant c^{(i)}\};$$

that is, set A denotes all types with higher equilibrium cost in the game with higher information, and set B denotes the rest of the types.

We also partition the routes of the network into two subsets $\mathcal{R}_A$ and $\mathcal{R}_B$, where

$$\mathcal{R}_A = \{r \in \mathcal{R}: \tilde{c}_r > c_r\}$$

and

$$\mathcal{R}_B = \{r \in \mathcal{R}: \tilde{c}_r \leqslant c_r\};$$

that is, $\mathcal{R}_A$ denotes all routes that have higher costs in the game with higher information, and $\mathcal{R}_B$ denotes the rest of the routes. We show the following claims.

Claim 1. *For any type $i \in A$ and any route $r \in \mathcal{R}_B$, we have $f_r^{(i)} = 0$; that is, for a given type i, if the equilibrium cost increases in the game with higher information, then the cost of all routes that type i was using (with strictly positive flow) also increases. This follows since if $r \notin \mathcal{R}_i$, then $f_r^{(i)} = 0$. Otherwise, $r \in \mathcal{R}_i$, which implies $r \in \tilde{\mathcal{R}}_i$ as well, where $\tilde{\mathcal{R}}_i$ denotes the set of available routes to type i in the expanded information set. Assuming $i \in A$ and $r \in \mathcal{R}_B$, we have*

$$c_r \geqslant \tilde{c}_r \geqslant \tilde{c}^{(i)} > c^{(i)},$$

where the first inequality follows from the definition of the set $\mathcal{R}_B$. The second inequality follows from the definition of ICWE. The third inequality follows from the definition of set A. The overall inequality and the definition of ICWE show that $f_r^{(i)} = 0$.

Claim 2. *For any type $i \in B$ and any route $r \in \mathcal{R}_A$, we have $\tilde{f}_r^{(i)} = 0$; that is, for a given route, if the cost of the route in the equilibrium increases in the game with higher information, then the equilibrium costs of all types that are using this route in the equilibrium of the higher information game also increases. This follows since if $r \notin \tilde{\mathcal{R}}_i$, then $\tilde{f}_r^{(i)} = 0$. Otherwise, $r \in \tilde{\mathcal{R}}_i$, which implies $r \in \mathcal{R}_i$, because $1 \notin B$ and the information set of all other types are fixed. Assuming $i \in B$ and $r \in \mathcal{R}_A$, we have*

$$\tilde{c}_r > c_r \geqslant c^{(i)} \geqslant \tilde{c}^{(i)},$$

where the first inequality follows from the definition of the set $\mathcal{R}_A$. The second inequality follows from the definition of ICWE. The third inequality follows from the definition of set B. The overall inequality and the definition of ICWE show that $\tilde{f}_r^{(i)} = 0$.

Claim 3. *Letting $s_A = \sum_{r \in \mathcal{R}_A} \sum_{i=1}^{K} f_r^{(i)}$, $\tilde{s}_A = \sum_{r \in \mathcal{R}_A} \sum_{i=1}^{K} \tilde{f}_r^{(i)}$, $s_B = \sum_{r \in \mathcal{R}_B} \sum_{i=1}^{K} f_r^{(i)}$, and $\tilde{s}_B = \sum_{r \in \mathcal{R}_B} \sum_{i=1}^{K} \tilde{f}_r^{(i)}$, we have $\tilde{s}_A \leqslant s_A$ and $\tilde{s}_B \geqslant s_B$.*

This follows from Claims 1 and 2. The traffic on the routes in $\mathcal{R}_A$ from $f^{(1:K)}$ is s_A, which is the entire traffic demand s_i for all $i \in A$ (Claim 1) and possibly some portion of the traffic demand s_j for $j \in B$. On the other hand, the traffic on the routes in $\mathcal{R}_A$ from $\tilde{f}^{(1:K)}$ is $\tilde{s}_A$, which contains only some portion of the traffic demand s_i for $i \in A$. Claim 2 implies that for all $j \in B$ the traffic demand $\tilde{s}_j$ is only sent on the routes in $\mathcal{R}_B$. This shows that $\tilde{s}_A \leqslant s_A$, which in turn leads to $\tilde{s}_B \geqslant s_B$ (see Figure 4 for an illustration of the partitioning and the flows).

Part (b) of Lemma 1 shows that there exists type i for which $\tilde{c}^{(i)} \leqslant c^{(i)}$, which in turn shows that set B is nonempty—also by the contradiction assumption $1 \in A$, which implies that both A and B are nonempty. Using Claim 1, if A is nonempty, then $\mathcal{R}_A \neq \emptyset$ as the flow $f^{(1:K)}$ of the types in A can only go to routes in $\mathcal{R}_A$. Also using Claim 2, since B is nonempty, we have $\mathcal{R}_B \neq \emptyset$ as the flow $\tilde{f}^{(1:K)}$ of the types in B can only go to routes in $\mathcal{R}_B$. Therefore, we have partitioned the routes of the network into two nonempty sets $\mathcal{R}_A$ and $\mathcal{R}_B$ such that $\tilde{c}_r > c_r$ for all $r \in \mathcal{R}_A$ and $\tilde{c}_r \leqslant c_r$ for all $r \in \mathcal{R}_B$. In other words, we have $\max_{r \in \mathcal{R}_A}\{c_r - \tilde{c}_r\} < 0$ and $\min_{r \in \mathcal{R}_B}\{c_r - \tilde{c}_r\} \geqslant 0$. We now have all the pieces to use Lemma 3, which yields

$$0 > \max_{r \in \mathcal{R}_A}\{c_r - \tilde{c}_r\} \geqslant \min_{r \in \mathcal{R}_B}\{c_r - \tilde{c}_r\} \geqslant 0,$$

which is a contradiction, completing the proof of part (a). □

Proof of Part (b). The proof of this part follows from Theorem 2. Let G be a non-SLI network. Using Theorem 2, one of the networks shown in Figure 3 must be embedded in G. Using Example 2, for all networks shown in Figure 3, there exists an assignment of cost functions and information sets for which IBP occurs.

To construct an example for G we start from the cost functions for which the embedded network features IBP (as shown in Example 2) and then following the steps of embedding, given in Definition 4, we will update the information sets as well as the cost functions in a way that IBP occurs in the final network G. The updates of information sets and cost functions are as follows.

(i) If the step of embedding is to divide an edge, we assign half of the original edge cost to each of the new edges and update the information set by adding both newly created edges to the same information set as of the original edge. This guarantees that the equilibrium flow of the network after dividing an edge is the same as the one before.

(ii) If the step of embedding is to add an edge, then we include that edge in none of the information sets (or equivalently assign cost infinity to it). This guarantees that the new edge is never used in any equilibrium.

(iii) If the step of embedding is to extend origin or destination, we let the cost of the new edge be $c(x)=x$ and update all of the information sets by adding this edge to them. Since this edge will be used by all types and the flow on it will not change, this step of embedding does not affect the equilibrium flow.

This construction establishes that since IBP is present in the initial network—that is, one of the networks shown in Figure 3—it will be present in the network G as well. This completes the proof of part (b).

Recall that in Remark 2 we showed for each of the networks shown in Figure 3 that there exist infinitely many cost functions for which IBP occurs. This shows that if IBP occurs in a network, then it occurs for infinitely many cost functions. This is because if IBP occurs in a network G, Theorem 3, part (a) implies G is not SLI, and Theorem 2 shows that one of the basic networks shown in Figure 3 is embedded in G. Finally, by construction of the proof of Theorem 3, part (b), the cost function configuration of the basic network can be extended to network G, showing that IBP occurs for infinitely many cost functions.

6.3. IBP with Restricted Information Sets

In this subsection, we show that restricting focus to networks with a much more specific information structure—whereby only one type does not know all the edges, and the change in question informs this type of all edges—allows us to establish that IBP does not occur in a larger set of networks. Interestingly, in this case, IBP does not occur in exactly the same set of networks on which BP does not occur, series-parallel networks, although the two concepts continue to be very different even under this more specific information structure. The similarity is that after the change, as in the classic Wardrop equilibrium setting studied for BP, there is no more heterogeneity among users. We first define IBP with restricted information sets and then state the characterization of network topology that leads to it.

Definition 8 (IBP with Restricted Information Sets). Consider a traffic network with multiple information types $(G,\mathscr{E}_{1:K},s_{1:K},\mathbf{c})$. IBP with restricted information sets occurs if there exist expanded information sets $\tilde{\mathscr{E}}_{1:K}$ with $\mathscr{E}_1\subset\tilde{\mathscr{E}}_1=\mathscr{E}$, and $\mathscr{E}_i=\tilde{\mathscr{E}}_i=\mathscr{E}$ for $i=2,\dots,K$, such that the equilibrium cost of type 1 in $(G,\tilde{\mathscr{E}}_{1:K},s_{1:K},\mathbf{c})$ is strictly larger than the equilibrium cost of type 1 in $(G,\mathscr{E}_{1:K},s_{1:K},\mathbf{c})$. We denote the equilibrium cost of type $i\in[K]$ before and after the expansion of information by $c^{(i)}$ and $\tilde{c}^{(i)}$, respectively.

Theorem 4. *IBP with restricted information sets does not occur if and only if the network G is SP. More specifically, we have the following*:

(a) *If G is SP, then for any network with multiple information sets $(G,\mathscr{E}_{1:K},s_{1:K},\mathbf{c})$ with arbitrary assignment of cost functions $\mathbf{c}$, K, traffic demands $s_{1:K}$, and information set $\mathscr{E}_1$, IBP with the restricted information sets does not occur.*

(b) *If G is not SP, then there exists an assignment of cost functions $\mathbf{c}$, K, traffic demands $s_{1:K}$, and information set $\mathscr{E}_1$ in which IBP with restricted information sets occurs.*

6.4. Extension to Multiple Origin–Destination Pairs

In this subsection, we consider networks with multiple information types and multiple origin–destination pairs as defined next.

Definition 9. Consider a graph $G=(V,\mathscr{E})$ containing m origin–destination pairs denoted by (O_i,D_i), $i\in[m]$. For any $i\in[m]$, there are K_i types of users, each with information set $\mathscr{E}_{i,j}\subseteq\mathscr{E}$, for $j\in[K_i]$. We refer to (i,j) as the type of a user, where $i\in[m]$ denotes the origin–destination pair of this type and $j\in[K_i]$ represents its information set. The traffic network with multiple information types and multiple origin–destination pairs is denoted by $(G,\{\mathscr{E}_{i,1:K_i}\}_{i=1}^m,\{s_{i,1:K_i}\}_{i=1}^m,\mathbf{c})$. We let $\mathscr{R}_{i,j}$ denote the set of routes available to a user of type (i,j) (i.e., routes formed by edges in $\mathscr{E}_{i,j}$). A feasible flow is a flow vector $f=(f^{(1,1:K_1)},\dots,f^{(m,1:K_m)})$ such that $f^{(i,1:K_i)}$ is a feasible flow for origin–destination pair (O_i,D_i).

We denote the total flow on an edge e by f_e, where $f_e=\sum_{i=1}^m\sum_{j=1}^{K_i}\sum_{r\in\mathscr{R}_{i,j}:e\in r}f_r^{(i,j)}$. Note that since G is an undirected graph, the total flow on each edge is the sum of the flows sent through that edge in either direction (see Lin et al. 2011, Holzman and Monderer 2015). The cost of a route r is defined as $c_r(f)=\sum_{e\in r}c_e(f_e)$. ICWE in this case is defined naturally as follows.

A feasible flow $f=(f^{(1,1:K_1)},\dots,f^{(m,1:K_m)})$ is an ICWE if for every $i\in[m]$ and $j\in[K_i]$ and every pair $r,\tilde{r}\in\mathscr{R}_{i,j}$ with $f_r^{(i,j)}>0$ we have

$$c_r(f)\leqslant c_{\tilde{r}}(f). \tag{3}$$

This implies that all routes of type (i,j) with positive flow have the same cost, which is smaller than or equal to the cost of any other route in $\mathscr{R}_{i,j}$. The equilibrium cost of type (i,j), denoted by $c^{(i,j)}$, is then given by the cost of any route in $\mathscr{R}_{i,j}$ with positive flow from type (i,j).

The existence of ICWE in this setting follows from an identical argument to that of Theorem 1. Finally, the definition of IBP for this extended setting is as follows.

Definition 10 (IBP with Multiple Origin–Destination Pairs). Consider a traffic network with multiple information types and multiple origin–destination pairs $(G, \{\mathscr{E}_{i,1:K_i}\}_{i=1}^m, \{s_{i,1:K_i}\}_{i=1}^m, \mathbf{c})$. IBP occurs if there exists an expanded information set $\{\tilde{\mathscr{E}}_{i,1:K_i}\}_{i=1}^m$ with $\mathscr{E}_{1,1} \subset \tilde{\mathscr{E}}_{1,1}$ and $\tilde{\mathscr{E}}_{i,j} = \mathscr{E}_{i,j}$ for all $(i,j) \neq (1,1)$, $i \in [m]$, $j \in [K_i]$, such that the equilibrium cost of type (1,1) in $(G, \{\tilde{\mathscr{E}}_{i,1:K_i}\}_{i=1}^m, \{s_{i,1:K_i}\}_{i=1}^m, \mathbf{c})$ is strictly larger than the equilibrium cost of type (1,1) in $(G, \{\mathscr{E}_{i,1:K_i}\}_{i=1}^m, \{s_{i,1:K_i}\}_{i=1}^m, \mathbf{c})$.

Note that the choice of type $(1,1)$ for information expansion is arbitrary and without loss of generality. We next establish a sufficient condition on the network topology under which IBP with multiple origin–destination pairs does not arise. We will use the following definitions from Chen et al. (2016).

Definition 11.

- For any origin–destination pair (O_i, D_i), the *relevant network i* denoted by $G_i = (V_i, \mathscr{E}_i)$ consists of all edges and nodes of G that belong to at least one route from O_i to D_i in G.
- For an SLI network G_i, each LI block has two terminal nodes, an origin and a destination, such that the origin is the first node and the destination is the last node in the block visited on any route in G_i. For two SLI networks G_i and G_j, a *coincident LI block* is a common LI block of G_i and G_j with the same set of terminal nodes, allowing the origin of one to be the destination of the other.

Note that the definition of relevant network G_i as well as its LI blocks depends only on the network G and the origin–destination pair (O_i, D_i), not on the information sets. From this definition, we next provide a sufficient condition for excluding IBP.

Proposition 4. *Let G be a graph with $m \geqslant 1$ origin–destination pairs. For any $i \in [m]$, let $G_i = (V_i, \mathscr{E}_i)$ be the relevant network for origin–destination pair (O_i, D_i). IBP does not occur if the following two conditions hold*:

(a) *For any $i \in [m]$, the network G_i is SLI.*

(b) *For any $i, i' \in [m]$ either $\mathscr{E}_i \cap \mathscr{E}_{i'} = \varnothing$ or $\mathscr{E}_i \cap \mathscr{E}_{i'}$ consists of all coincident blocks of G_i and $G_{i'}$.*

Proof. We let f and $\tilde{f}$ denote the equilibrium flows before and after the expansion of the information set of type $(1,1)$. To reach a contradiction, suppose that $\tilde{c}^{(1,1)} > c^{(1,1)}$. Using part (b) of Lemma 2 and condition (a) of the proposition, the equilibrium cost of type $(1,1)$ users is the sum of the equilibrium cost of the LI blocks of G_1. Since $\tilde{c}^{(1,1)} > c^{(1,1)}$, there exists an LI block of G_1 for which the equilibrium cost after expanding information set of type $(1,1)$ increases. We denote this LI block by G^* and its corresponding origin and destination by O^* and D^*, respectively. Using condition (b) for any $i \neq 1$, we have one of the following two cases: (i) G_i does not have any common edge with G^*, and therefore none of the route flows of (O_i, D_i) goes through any edge of G^*; or (ii) O^* and D^* belong to all routes of G_i, and therefore all route flows of (O_i, D_i) go through G^*. We let C be the set of indices of such origin–destination pairs (i.e., $C = \{i \in [m]\colon O^*, D^* \in r, \forall\, r \in G_i\}$). We next define a traffic network with single origin–destination pair (O^*, D^*) over G^* for which IBP has occurred. The types of users are $(\bigcup_{i \in C}\{(i,j)\colon j \in [K_i]\}) \cup \{(1,j)\colon j \in [K_1]\}$ with their corresponding traffic demands. Note that for all $i \in C$, even though our definition of coincident LI block allows the route flows of (O_i, D_i) to go from O^* to D^* in either direction, without loss of generality, we can assume that route flows go from O^* to D^*. This is because the cost of any edge is a function of the sum of the flows that passes through that edge in either direction, and reversing the flows does not change the equilibrium flows on edges. Using Lemma 2, part (a), the restriction of equilibrium flows f and $\tilde{f}$ to G^* are equilibrium flows for the congestion game with multiple information types and a single origin–destination pair defined on G^*. Note that G^* is an LI network, and the equilibrium cost of type $(1,1)$ users after expanding their information set has gone up, which is a contradiction using Theorem 3. □

Figure 5(a) shows two SLI networks with their corresponding LI blocks, and Figure 5(b) shows a graph with two origin–destination pairs, which satisfies our sufficient condition.

The next example shows that the conditions of Proposition 4 are not necessary for nonoccurrence of IBP.

Example 3. Consider the network G shown in Figure 6. The common LI block of relevant networks G_1 and G_2 is G itself, which is not a coincident LI block because the sets of terminals of this block for G_1 and G_2 are different. Therefore, this network does not satisfy the conditions of Proposition 4. However, in Appendix A.4.4 we show that for any set of edge cost functions, IBP does not occur in this network.

In concluding this subsection we should note that BP with multiple origin–destination pairs has been studied in Epstein et al. (2009), Lin et al. (2011), Fujishige et al. (2017), Holzman and Monderer (2015), and Chen et al. (2016). In particular, Chen et al. (2016) provide a full characterization of network topologies for which BP occurs with multiple origin–destination pairs. BP as defined in Chen et al. (2016) occurs if adding an edge (decreasing cost of an edge) increases the equilibrium cost of the users of one of the origin–destination pairs, even if that edge is never used by the users of that origin–destination pair. With this definition it is possible to have a network for which IBP does not occur while BP occurs. For instance, BP occurs in the network considered in Example 3 (see Chen et al. 2016), while we showed IBP does not occur in this network.

Figure 5. (Color online) (a) Two SLI Networks with Their Corresponding LI Blocks; (b) A Graph with Two OD Pairs for Which IBP Does Not Occur

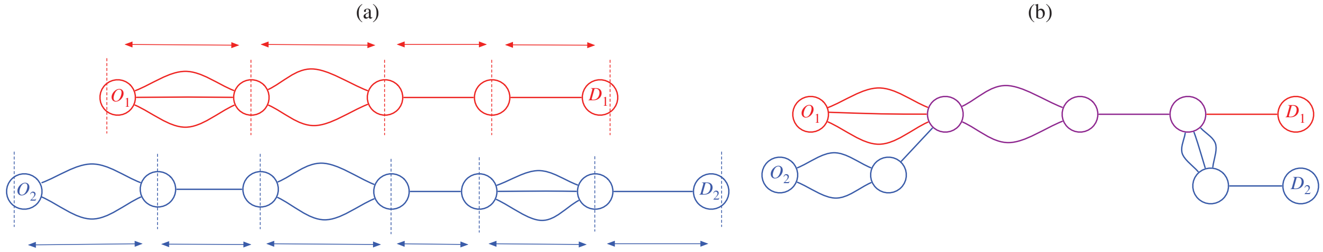

Figure 6. Example 3: IBP Does Not Occur on This Network

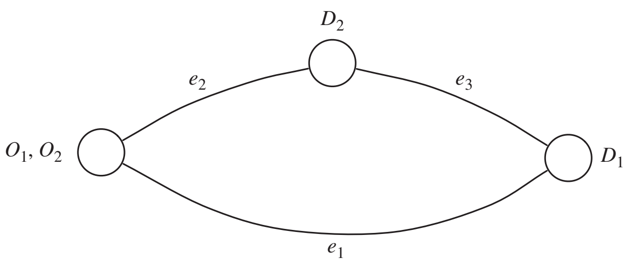

7. Efficiency of Information Constrained Wardrop Equilibrium

In this section, we provide bounds on the inefficiency of ICWE. We show that the worst-case inefficiency remains the same as the standard Wardrop equilibrium, even though our notion of ICWE is considerably more general than the Wardrop equilibrium since it allows for a rich amount of heterogeneity among users.

We start by defining the social optimum defined as the feasible flow vector that minimizes the total cost over all edges. We focus on aggregate efficiency loss defined as the ratio of total cost experienced by all users at social optimum and ICWE. We provide tight bounds on this measure of efficiency loss that are realized for different classes of cost functions. We also consider type-specific efficiency loss defined as the ratio of total cost experienced by type i users at social optimum and ICWE. We show that the bounds in this case are different from the ones in the standard Wardrop equilibrium.

Given a traffic network with multiple information types $(G, \mathscr{E}_{1:K}, s_{1:K}, \mathbf{c})$, we define the social optimum, denoted by $f_{\text{so}}^{(1:K)} = (f_{\text{so}}^{(1)}, \ldots, f_{\text{so}}^{(K)})$ (or simply f_{so}), as the optimal solution of the following optimization problem:

$$\min \sum_{e\in\mathscr{E}} f_e\, c_e(f_e),$$

$$f_e = \sum_{i=1}^{K} \sum_{r\in\mathscr{R}_i : e\in r} f_r^{(i)},$$

$$\sum_{r\in\mathscr{R}_i} f_r^{(i)} = s_i, \quad \text{and} \quad f_r^{(i)} \geqslant 0 \quad \text{for all } r\in\mathscr{R}_i \text{ and } i. \qquad (4)$$

This optimization problem minimizes the total cost over all edges incurred by all users of all types. Under the assumption that each cost function is continuous, it follows that the optimal solution of problem (4) and hence a social optimum always exists. We denote the total cost of a feasible flow $f^{(1:K)}$ by

$$C(f^{(1:K)}) \triangleq \sum_{e\in\mathscr{E}} f_e c_e(f_e).$$

Similarly, for a feasible flow $f^{(1:K)}$, we define the total cost incurred by type i users as

$$C^{(i)}(f^{(1:K)}) \triangleq \sum_{e\in\mathscr{E}} f_e^{(i)} c_e(f_e).$$

Consequently, we define the socially optimal cost of type i as $C_{\text{so}}^{(i)} = C^{(i)}(f_{\text{so}}^{(1:K)})$ for $i\in[K]$ and the overall cost (over all types) of social optimum as $C_{\text{so}} = C(f_{\text{so}}^{(1:K)})$. Similarly, we define the equilibrium cost of type i as $C_{\text{cwe}}^{(i)} = C^{(i)}(f_{\text{cwe}}^{(1:K)})$ for $i\in[K]$ and the overall cost (over all types) of ICWE as $C_{\text{cwe}} = C(f_{\text{cwe}}^{(1:K)})$, where $f_{\text{cwe}}^{(1:K)}$ (or simply f_{cwe}) denotes an ICWE. Note that $C^{(i)}(f_{\text{cwe}}^{(1:K)})$ is different from the equilibrium cost of type i denoted by $c^{(i)}$, as the latter notion is the cost per unit of flow and the former is the aggregate cost. The relation between these two is simply $C_{\text{cwe}}^{(i)} = s_i\, c^{(i)}$, $i\in[K]$.

The following result from Roughgarden and Tardos (2002) and Correa et al. (2005) presents bounds on the efficiency loss of the Wardrop equilibrium, which provides bounds on the efficiency loss of ICWE in a traffic network with a single information type, with $\mathscr{E}_1 = \mathscr{E}$ denoted by $(G, \mathscr{E}_1, s_1, \mathbf{c})$.

Proposition 5 (Roughgarden and Tardos (2002)). *Consider a traffic network with a single information type $(G, \mathscr{E}_1, s_1, \mathbf{c})$. Let f_{we} be a Wardrop equilibrium, and let f_{so} be a social optimum. Then, we have*

(a) $\inf_{(G,\mathscr{E}_1,s_1,\mathbf{c}):\, c_e \text{convex}} C_{\text{so}}/C_{\text{we}} = 0.$

(b) *Suppose $c_e(x)$ is an affine function for all $e\in\mathscr{E}$. Then, we have $C_{\text{so}}/C_{\text{we}} \geqslant \frac{3}{4}$, and this bound is tight.*

(c) *Let $\mathscr{C}$ be a class of latency functions, and let $\beta(\mathscr{C}) = \sup_{c\in\mathscr{C},\, x\geqslant 0} \beta(c,x)$, where*

$$\beta(c,x) = \max_{z\geqslant 0} \frac{z\,(c(x)-c(z))}{x\, c(x)}.$$

Then, we have $C_{\text{so}}/C_{\text{we}} \geqslant 1-\beta(\mathscr{C})$, and the bound is tight.

Our next result shows that Proposition 5 holds exactly for ICWE, indicating that within the class of heterogeneous, information-constrained traffic equilibria we consider, the worst-case scenario occurs for networks with homogeneous users.

Proposition 6. *Consider a traffic network with multiple information types* $(G, \mathscr{E}_{1:K}, s_{1:K}, \mathbf{c})$. *Let* f_{cwe} *be an ICWE, and let* f_{so} *be a social optimum. Then, we have*

(a) $\inf_{(G,\mathscr{E}_{1:K},s_{1:K},\mathbf{c}):c_e \text{ convex}} C_{\text{so}}/C_{\text{cwe}} = 0$.

(b) *Suppose* $c_e(x)$ *is an affine function for all* $e \in \mathscr{E}$. *Then, we have* $C_{\text{so}}/C_{\text{cwe}} \geqslant \frac{3}{4}$, *and this bound is tight.*

(c) *Let* $\mathscr{C}$ *be a class of latency functions, and let* $\beta(\mathscr{C}) = \sup_{c\in\mathscr{C}, x\geqslant 0} \beta(c,x)$, *where*

$$\beta(c,x) = \max_{z\geqslant 0} \frac{z\,(c(x)-c(z))}{x\,c(x)}.$$

Then, we have $C_{\text{so}}/C_{\text{cwe}} \geqslant 1-\beta(\mathscr{C})$, *and the bound is tight.*

Proof. We first show that for any type i, and any feasible flow $f^{(i)}$ for this type, we have

$$\sum_{e\in\mathscr{E}} c_e(f_{e,\text{cwe}})(f^{(i)}_{e,\text{cwe}} - f^{(i)}_e) \leqslant 0. \tag{5}$$

The reason is that in ICWE each type uses only the routes with the minimal costs. Therefore, for any type i and any feasible flow $f^{(i)}$ for type i, we have

$$\sum_{r\in\mathscr{R}_i} c_r(f^{(1:K)}_{\text{cwe}}) f^{(i)}_{r,\text{cwe}} \leqslant \sum_{r\in\mathscr{R}_i} c_r(f^{(1:K)}_{\text{cwe}}) f^{(i)}_r.$$

This leads to

$$\begin{aligned} 0 &\geqslant \sum_{r\in\mathscr{R}_i} c_r(f^{(1:K)}_{\text{cwe}})(f^{(i)}_{r,\text{cwe}} - f^{(i)}_r) \\ &= \sum_{r\in\mathscr{R}_i}\left(\sum_{e:e\in r} c_e(f_{e,\text{cwe}})\right)(f^{(i)}_{r,\text{cwe}} - f^{(i)}_r) \\ &= \sum_{e\in\mathscr{E}} c_e(f_{e,\text{cwe}}) \sum_{r\in\mathscr{R}_i:e\in r} (f^{(i)}_{r,\text{cwe}} - f^{(i)}_r) \\ &= \sum_{e\in\mathscr{E}} c_e(f_{e,\text{cwe}})(f^{(i)}_{e,\text{cwe}} - f^{(i)}_e), \end{aligned}$$

which is the desired inequality, showing Equation (5). We next proceed with the proof.

Part (a). This holds because a traffic network with one type is a special case of traffic network with multiple information types, and part (a) of Proposition 5 shows that the infimum is zero.

Part (b). Using Equation (5) for $f^{(i)} = f^{(i)}_{\text{so}}$ for any $i \in [K]$, and taking summation over all types $i \in [K]$, we obtain

$$\begin{aligned} C_{\text{cwe}} &= \sum_{e\in\mathscr{E}} f_{e,\text{cwe}} c_e(f_{e,\text{cwe}}) \\ &= \sum_{i=1}^K \sum_{e\in\mathscr{E}} c_e(f_{e,\text{cwe}}) f^{(i)}_{e,\text{cwe}} \leqslant \sum_{i=1}^K \sum_{e\in\mathscr{E}} c_e(f_{e,\text{cwe}}) f^{(i)}_{e,\text{so}} \\ &= \sum_{e\in\mathscr{E}} c_e(f_{e,\text{cwe}}) \sum_{i=1}^K f^{(i)}_{e,\text{so}} = \sum_{e\in\mathscr{E}} f_{e,\text{so}} c_e(f_{e,\text{cwe}}) \\ &= \sum_{e\in\mathscr{E}} f_{e,\text{so}} c_e(f_{e,\text{so}}) + \sum_{e\in\mathscr{E}} f_{e,\text{so}}(c_e(f_{e,\text{cwe}}) - c_e(f_{e,\text{so}})) \\ &\leqslant \sum_{e\in\mathscr{E}} f_{e,\text{so}} c_e(f_{e,\text{so}}) + \tfrac{1}{4}\sum_{e\in\mathscr{E}} f_{e,\text{cwe}} c_e(f_{e,\text{cwe}}), \end{aligned}$$

where the last inequality comes from the fact that with $c_e(x) = a_e x + b_e$ for $b_e, a_e \geqslant 0$, we have

$$\begin{aligned} &f_{e,\text{so}}(c_e(f_{e,\text{cwe}}) - c_e(f_{e,\text{so}})) \\ &\quad = a_e f_{e,\text{so}}(f_{e,\text{cwe}} - f_{e,\text{so}}) \leqslant \tfrac{1}{4} f^2_{e,\text{cwe}} a_e \leqslant \tfrac{1}{4} f_{e,\text{cwe}} c_e(f_{e,\text{cwe}}). \end{aligned}$$

The proof of tightness follows from part (b) of Proposition 5 as a traffic network with one type is a special case of a traffic network with multiple information types.

Part (c). Using the same argument as in part (b), we obtain

$$\begin{aligned} C_{\text{cwe}} &= \sum_{e\in\mathscr{E}} f_{e,\text{cwe}} c_e(f_{e,\text{cwe}}) \\ &\leqslant \sum_{e\in\mathscr{E}} f_{e,\text{so}} c_e(f_{e,\text{so}}) + \sum_{e\in\mathscr{E}} f_{e,\text{so}}(c_e(f_{e,\text{cwe}}) - c_e(f_{e,\text{so}})) \\ &\leqslant \sum_{e\in\mathscr{E}} f_{e,\text{so}} c_e(f_{e,\text{so}}) + \beta(\mathscr{C}) \sum_{e\in\mathscr{E}} f_{e,\text{cwe}} c_e(f_{e,\text{cwe}}), \end{aligned}$$

where the last inequality comes from the fact that

$$\begin{aligned} f_{e,\text{so}}(c_e(f_{e,\text{cwe}}) - c_e(f_{e,\text{so}})) &\leqslant \beta(c_e, f_{e,\text{cwe}}) f_{e,\text{cwe}} c_e(f_{e,\text{cwe}}) \\ &\leqslant \beta(\mathscr{C}) f_{e,\text{cwe}} c_e(f_{e,\text{cwe}}). \end{aligned}$$

The proof of the tightness follows from part (c) of Proposition 5. □

In concluding this section, we should note that in this environment with heterogeneous users, there are alternatives to our formulation of the social optimum problem, which considers the "utilitarian" social optimum, summing over the costs of all groups. An alternative would be to consider a weighted sum or focus on the class of users suffering the greatest costs. We next illustrate that if we focus on type-specific costs, even with affine cost functions, some groups of users may have worse than 3/4 performance relative to the social optimum.

Example 4. Consider the network shown in Figure 7 with $\mathscr{E}_1 = \{e_1\}$, $\mathscr{E}_2 = \{e_1, e_2\}$. The ICWE is $f^{(1)}_{e_1,\text{cwe}} = s_1$ and $f^{(2)}_{e_1,\text{cwe}} = 1/a - s_1$, $f^{(2)}_{e_2,\text{cwe}} = s_2 - 1/a + s_1$. The equilibrium costs are $C^{(1)}_{\text{cwe}} = s_1$ and $C^{(2)}_{\text{cwe}} = s_2$. The social optimum is $f^{(1)}_{e_1,\text{so}} = s_1$ and $f^{(2)}_{e_1,\text{so}} = 1/(2a) - s_1$, $f^{(2)}_{e_2,\text{so}} = s_2 - 1/(2a) + s_1$. The corresponding costs are $C^{(1)}_{\text{so}} = s_1/2$ and $C^{(2)}_{\text{so}} = -1/(4a) + s_2 + s_1/2$ (assuming $1/(2a) \geqslant s_1$ and $s_2 \geqslant 1/a - s_1$). Therefore, we have

$$\frac{C^{(1)}_{\text{so}}}{C^{(1)}_{\text{cwe}}} = \frac{1}{2}, \quad \frac{C^{(2)}_{\text{so}}}{C^{(2)}_{\text{cwe}}} = \frac{-1/(4a) + s_2 + s_1/2}{s_2}, \quad \text{and}$$

$$\frac{C^{(1)}_{\text{so}} + C^{(2)}_{\text{so}}}{C^{(1)}_{\text{cwe}} + C^{(2)}_{\text{cwe}}} = \frac{s_1 + s_2 - 1/(4a)}{s_1 + s_2}.$$

Figure 7. Example 4: Type-Specific Efficiency Loss vs. Aggregate Efficiency Loss

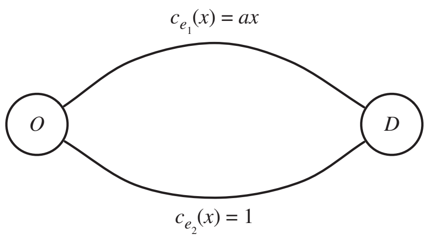

We next show that the ratio of the aggregate costs is greater than or equal to 3/4. We have $s_1 + s_2 \geqslant 1/a$, which leads to

$$\frac{C_{\text{so}}}{C_{\text{cwe}}} = \frac{s_1 + s_2 - 1/(4a)}{s_1 + s_2} = 1 - \frac{1}{4a}\frac{1}{s_1 + s_2} \geqslant 1 - \frac{1}{4a}a = \frac{3}{4}.$$

However, the type-specific efficiency loss can be smaller than $\frac{3}{4}$ as we have $(C^{(1)}_{\text{so}}/C^{(1)}_{\text{cwe}}) < \frac{3}{4}$.

8. Concluding Remarks

GPS-based route guidance systems, such as Waze or Google Maps, are rapidly spreading among drivers because of their promise of reduced delays as they inform their users about routes that they were not aware of or help them choose dynamically between routes depending on recent levels of congestion. Nevertheless, there is no systematic analysis of the implications for traffic equilibria of additional information provided to subsets of users. In this paper, we systematically studied this question. We first extended the class of standard congestion games used for analysis of traffic equilibria to a setting where users are heterogeneous because of their different information sets about available routes. In particular, each user's information set contains information about a subset of the edges in the entire road network, and drivers can only utilize routes consisting of edges that are in their information sets. We defined the notion of ICWE, an extension of the classic Wardrop equilibrium notion, and established the existence and essential uniqueness of ICWE.

We then turned to our main focus, which we formulate in the form of IBP. IBP asks whether users receiving additional information can become worse off. Our main result is a comprehensive answer to this question. We showed that in any network in the SLI class, which is a strict subset of series-parallel network, IBP cannot occur, and in any network that is not in the SLI class, there exists a configuration of edge-specific cost functions for which IBP will occur. The SLI class is made up of networks that join linearly independent networks in series, and linearly independent networks are those for which every path between the origin and destination contains at least one edge that is not in any other such path. This is the property that enables us to prove that IBP cannot occur in any SLI network. We also showed that any network that is not in the SLI class necessarily embeds at least one of a specific set of basic networks, and then we used this property to show that IBP will occur for some cost configurations in any non-SLI network. We further proved that whether a given network is SLI can be determined in linear time. Finally, we also established that the worst-case inefficiency performance of ICWE is no worse than the standard Wardrop equilibrium with one type of user.

There are several natural research directions that are opened up by our study. These include the following:

- Our analysis focused on the effect of additional information on the set of users receiving the information; for what classes of networks is additional information very harmful for other users? This question is important from the viewpoint of fairness and other social objectives. We may like that users utilizing route guidance systems are experiencing lower delays but not if this comes at the cost of significantly longer delays for others.
- How "likely" are the cost function configurations that cause IBP to occur in non-SLI networks? This question is important for determining, ex ante before knowing the exact traffic flows, whether additional information for some sets of users, coming, for example, from route guidance systems, might be harmful.
- Is there an "optimal information" configuration for users of a traffic network? Specifically, one could consider the following question: Given the traffic demands of K types, $s_1, \ldots, s_K$, find the information sets $\mathscr{E}_1, \ldots, \mathscr{E}_K$ that generate the minimum overall cost for all types in an ICWE. This question is related to Roughgarden (2001, 2006), who investigate the question of finding the subnetwork of the initial network that leads to an optimal equilibrium cost with one type of user.
- We established a sufficient condition under which IBP does not occur on a traffic network with multiple origin–destination pairs. One natural question is to find a sufficient and necessary condition for this problem.
- Finally, our study poses an obvious empirical question, complementary to similar studies for the Braess' paradox: Are there real-world settings where we can detect IBP?

Acknowledgments

The authors would like to thank anonymous referees and editors for their useful comments and suggestions. The authors also thank participants at several seminars and conferences for useful suggestions and comments. In particular, they thank Nicolas Stier, Rakesh Vohra, Ozan Candogan, and Saeed Alaei for helpful comments.

Appendix A

A.1. Proofs of Section 3

A.1.1. Proof of Proposition 1. Since for any $e \in \mathscr{E}$ the function $c_e(\cdot)$ is nondecreasing, $\int_0^{f_e} c_e(z)\,dz$ as a function of $f_r^{(i)}$ is convex and continuously differentiable.

Claim 1. *If $f^{(1:K)}$ is an optimal solution of* (2), *then it is an ICWE.*

Since the objective function is convex and the constraints are affine functions, regularity conditions hold and Karush-Kuhn-Tucker (KKT) conditions are satisfied; that is, there exists $\mu_{i,r} \leqslant 0$ and λ_i such that for all $i \in [K]$ and $r \in \mathcal{R}_i$ we have

$$\frac{\partial}{\partial f_r^{(i)}}\left(\sum_{e\in\mathcal{E}}\int_0^{f_e} c_e(z)\,dz - \sum_{i=1}^{K}\lambda_i\left(\sum_{r\in\mathcal{R}_i} f_r^{(i)} - s_i\right) + \sum_{r,i}\mu_{r,i} f_r^{(i)}\right) = 0, \quad \text{(A.1)}$$

where $\mu_{r,i} = 0$ for $f_r^{(i)} > 0$ (Bertsekas 1999, chap. 3). We show that the flow $f^{(1:K)}$ is an ICWE with the equilibrium cost of type i being λ_i. First, note that $f^{(1:K)}$ is a feasible flow by the constraints of (2). Second, we can rewrite (A.1) as

$$\sum_{e\in\mathcal{E}}\frac{\partial f_e}{\partial f_r^{(i)}} c_e(f_e) = \sum_{e\in\mathcal{E}:\, e\in r} c_e(f_e) = \begin{cases} = \lambda_i & \text{if } f_r^{(i)} > 0,\\ \geqslant \lambda_i & \text{if } f_r^{(i)} = 0,\end{cases} \quad \text{(A.2)}$$

where we used $\mu_{r,i} = 0$ for $f_r^{(i)} > 0$ in the first case and $\mu_{r,i} \leqslant 0$ for $f_r^{(i)} = 0$ in the second case. This is exactly the definition of ICWE, which completes the proof of Claim 1.

Claim 2. *If $f^{(1:K)}$ is an ICWE, then it is an optimal solution of* (2).

We let the equilibrium cost of type i users be λ_i, which leads to the following relation:

$$\sum_{e\in\mathcal{E}:\, e\in r} c_e(f_e) = \begin{cases} = \lambda_i & \text{if } f_r^{(i)} > 0,\\ \geqslant \lambda_i & \text{if } f_r^{(i)} = 0.\end{cases} \quad \text{(A.3)}$$

For all $i \in [K]$ and $r \in \mathcal{R}_i$, if $f_r^{(i)} > 0$, then we define $\mu_{i,r} = 0$, and if $f_r^{(i)} = 0$, then we define $\mu_{i,r} = \lambda_i - \sum_{e\in\mathcal{E}:e\in r} c_e(f_e)$. First, note that $\mu_{i,r} \leqslant 0$, and if $f_r^{(i)} > 0$, then $\mu_{i,r} = 0$. Second, note that

$$\frac{\partial}{\partial f_r^{(i)}}\left(\sum_{e\in\mathcal{E}}\int_0^{f_e} c_e(z)\,dz - \sum_{i=1}^{K}\lambda_i\left(\sum_{r\in\mathcal{R}_i} f_r^{(i)} - s_i\right) + \sum_{r,i}\mu_{r,i} f_r^{(i)}\right) = 0. \quad \text{(A.4)}$$

Therefore, the flow $f^{(1:K)}$, together with λ_i and $\mu_{i,r}$, satisfies the KKT conditions. Since the objective function of (2) is convex and the constraints are affine functions, KKT conditions are sufficient for optimality (Bertsekas 1999, chap. 3), proving the claim. □

A.1.2. Proof of Theorem 1. The set of feasible flows $f^{(1:K)}$ is a compact subset of a $K|\mathcal{R}|$-dimensional Euclidean space. Since edge cost functions are continuous, the potential function is also continuous. The Weierstrass extreme value theorem establishes that optimization problem (2) attains its minimum, which, by Proposition 1, is an ICWE.

We next show that in two different equilibria $f^{(1:K)}$ and $\tilde{f}^{(1:K)}$, the equilibrium cost for each type is the same. By Proposition 1, both $f^{(1:K)}$ and $\tilde{f}^{(1:K)}$ are optimal solutions of (2). Since $\Phi(\cdot)$ is a convex function, we have

$$\Phi(\alpha f^{(1:K)} + (1-\alpha)\tilde{f}^{(1:K)}) \leqslant \alpha\Phi(f^{(1:K)}) + (1-\alpha)\Phi(\tilde{f}^{(1:K)})$$

for any $\alpha \in [0,1]$. Since $\Phi(f^{(1:K)})$ and $\Phi(\tilde{f}^{(1:K)})$ are both equal to the optimal value of (2), and for each e, the function $\int_0^{f_e} c_e(z)\,dz$ is convex (its derivative with respect to f_e is $c_e(f_e)$, which is nondecreasing), the functions $\int_0^{f_e} c_e(z)\,dz$ for any $e \in \mathcal{E}$ must be linear between values of f_e and $\tilde{f}_e$. This shows that all cost functions c_e are constant between f_e and $\tilde{f}_e$, and in particular, the equilibrium costs are the same. □

A.2. Proofs of Section 4

A.2.1. Proof of Equivalence in Definition 5. We first show that each LI network G is the result of attaching several LI blocks in series. This follows by induction on the number of edges. Using Definition 3, G is either the result of attaching two LI networks in parallel or the result of attaching an LI network and a single edge in a series. If G is the result of attaching two LI networks in parallel, then G is biconnected and so is an LI block. If G is the result of attaching an LI network G_1 with a single edge, then the single edge is an LI block and, by the induction hypothesis, G_1 is a series of several LI blocks. Therefore, G is the result of attaching several LI blocks in a series.

We next show that the following two definitions are equivalent.

- An SLI network is either a single LI network or the connection of two SLI networks in series. We let SET1 denote the set of such networks.
- An SLI network consists of attaching several LI blocks in series. We let SET2 denote the set of such networks.

We show that SET1 = SET2 by induction on the number of edges; that is, we suppose that for any network with its number of edges less than or equal to m that these two sets are equal and then show that for networks with $m+1$ edges that the two sets are equal as well (note that the base of this induction for $m = 1$ corresponds to a single edge, which evidently holds).

- If a network G belongs to SET1, then either it is a single LI network or it is the result of attaching two SLI networks in series. In the former case, it belongs to SET2 as we have shown each LI network is the result of attaching several LI blocks. In the latter case, by the induction hypothesis, both SLI subnetworks are the series of several LI blocks and so is their attachment in series. This shows SET1 ⊆ SET2.
- If a network G belongs to SET2, then either it is a single LI block or it is the result of attaching several LI blocks in series. In the former case, by definition, it belongs to SET1. In the latter case, we let G_1 denote the LI block that contains origin and the series of the rest of LI blocks by G_2. By definition, G_1 is SLI as it is a single LI block and G_2 is SLI by the induction hypothesis. Therefore, the series attachment of G_1 and G_2 belongs to SET1. This shows SET2 ⊆ SET1, completing the proof.

A.2.2. Proof of Theorem 2. We first show that if a network G belongs to the class SLI, then none of the networks shown in Figure 3 is embedded in it. First note that since all networks in the class SLI are series parallel, using part (b) of Proposition 2 implies that the Wheatstone network shown in Figure 3(a) is not embedded in it. The SLI network G consists of several LI blocks that are attached in a series. Using part (a) of Proposition 2, none of the networks shown in Figure 3 (i.e., networks shown in Figures 3(b)–3(i)) can be embedded in one of the LI blocks.

We next show that they cannot be embedded in the series of two LI blocks as well. We let G_1 and G_2 be two LI blocks that are attached in series where the resulting network from this attachment is H. Also, we let the node c be the attaching node of these two networks. We will show that the network shown in Figure 3(b) cannot be embedded in H (a similar argument shows that the rest of the networks shown in Figure 3 cannot be embedded in it). To reach to a contradiction,

Figure A.1. (Color online) Proof of Theorem 2: G_1 Is Not LI and G_2 Has at Least One Route from O to D

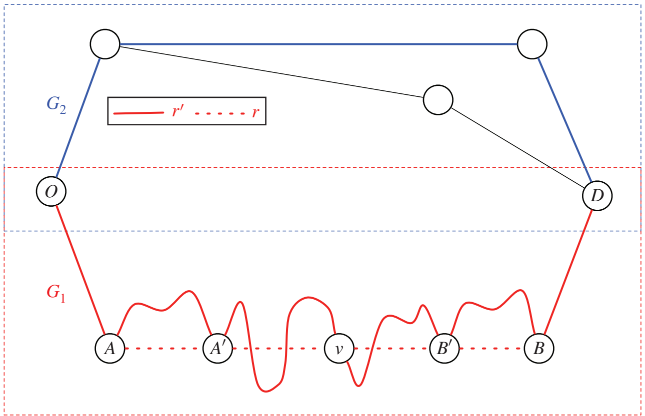

we suppose the contrary—that is, H is obtained from the network shown in Figure 3(b) by applying the embedding procedure described in Definition 4. We define the corresponding routes to e_5, e_1e_4, and e_2e_3 in H by r_3, r_1, and r_2. Formally, we start from $r_3 = e_5$, $r_1 = e_1e_4$, and $r_2 = e_2e_3$ in the network shown in Figure 3(b), and at each step of the embedding procedure whenever we divide an edge on r_i $(i = 1, 2, 3)$, we will update r_i by adding that edge, and whenever we extend the origin or destination, we will add the new edge to all r_i's. Given this construction, in the network H we have three routes, r_3, r_1, and r_2, where r_1 and r_2 have a common node and do not have any common node (except O and D) with r_3. This is a contradiction, as all routes in H must have node c in common. This completes the proof of the first part.

We next show that if none of the networks shown in Figure 3 is embedded in G, then G belongs to the class SLI. Proposition 2(b) implies that since Figure 3(a) is not embedded in G, it is series parallel. We next show that given a series-parallel network G, if G is not SLI, then we can find an embedding of one of the networks shown in Figures 3(b)–3(i) in it. The proof is by induction on the number of edges of G. Following Definition 2, consider the last building step of the network G. If the last step is attaching two networks G_1 and G_2 in a series, then assuming that G is not SLI, we conclude that either G_1 or G_2 is not SLI (or neither are). Therefore, by the induction hypothesis, we can find an embedding of one of the networks shown in Figures 3(b)–3(i) in either G_1 or G_2, which in turn shows that it is embedded in G. If the last step is attaching two networks G_1 and G_2 in parallel, then it must be the case that either G_1 or G_2 is not LI. This is because, otherwise, the parallel attachment of two LI networks is LI (Definition 3) and hence SLI, which contradicts the fact that G is not SLI. Without loss of generality, we let the network that is not LI be G_1. Therefore, part (a) of Proposition 2 shows that there exist two routes r and r' and a vertex v common to both routes such that both sections r_{Ov} and r'_{Ov} as well as r_{vD} and r'_{vD} are not equal (note that $v \notin \{O, D\}$ because otherwise, if $v = O$, then $r_{Ov} = r'_{Ov}$, as both are the single node O).

Note that using part (b) of Proposition 2, there is a way to index vertices such that along any route, the vertices have increasing indices. We let A be the last vertex (with the prescribed indexing) before which the two routes r and r' become the same (this vertex can be O itself). Since v is the common vertex of these two routes and $r_{Ov} \neq r'_{Ov}$, such a vertex exists. Because v is a common vertex of r and r', the two routes r and r' have a common vertex between A and v. We let A' be the first such vertex (it can be v itself). Similarly, we define B as the first vertex after which r and r' become the same (B can be D itself) and B' as the last vertex after v for which r and r' coincide (B' can be v itself). Given these definitions for the nodes v, A, A', B, and B', we know that $r_{AA'}$ (the path between A and A' on r) and $r'_{AA'}$ (the path between A and A' on r') do not have any vertex in common, and similarly, $r_{BB'}$ and $r'_{BB'}$ do not have any vertex in common. The definition of the nodes A, A', B, and B' is illustrated in Figure A.1. Next, we show that one of the networks shown in Figures 3(b)–3(i) is embedded in G. We have the following cases:

- $A = O$, $B = D$, $A' = v$, and $B' = v$: In this case, the network shown in Figure 3(b) is embedded in G. This is because there are two disjoint paths from O to v and from v to D, and there is at least one path from O to D in G_2. Since any other edge and vertex of the network belong to a path that connects O to D, we can construct the graph G by starting from the network shown in Figure 3(b) and applying the embedding procedure.
- $A = O$, $B = D$, and $A' \neq v$ or $B' \neq v$: In this case, the network shown in Figure 3(c) is embedded in G. This is because there is at least one path from O to D in G_2, and the network shown in Figure 3(c) is embedded in G_1. To see this, note that the edges e_1 and e_2 are embedded in the section of the routes r and r' between O and A', and the edges e_3 and e_4 are embedded in the section of the routes r and r' between B' and D. Also, note that the single edge e_6 is embedded in the network between A' and B' (the single edge is embedded in any network).
- $A \neq O$, $B = D$, and $A' = v$ and $B' = v$: The network shown in Figure 3(d) is embedded in G.
- $A = O$, $B \neq D$, and $A' = v$ and $B' = v$: The network shown in Figure 3(e) is embedded in G.
- $A = O$, $B \neq D$, and $A' \neq v$ or $B' \neq v$: The network shown in Figure 3(f) is embedded in G.
- $A \neq O$, $B = D$, and $A' \neq v$ or $B' \neq v$: The network shown in Figure 3(g) is embedded in G.
- $A \neq O$, $B \neq D$, and $A' \neq v$ or $B' \neq v$: The network shown in Figure 3(h) is embedded in G.
- $A \neq O$, $B \neq D$, and $A' = v$ and $B' = v$: The network shown in Figure 3(i) is embedded in G.

This completes the proof.

A.2.3. Proof of Proposition 3. We use the following results and definitions in this proof.

Proposition 7 (Valdes et al. 1979). *A network is series parallel if following the steps S and P shown in Figure A.2 in any order turns the network into a single edge connecting the origin to the destination. Moreover, if a network is series parallel, then in linear time* $O(|\mathscr{E}| + |V|)$ *we can obtain a binary tree decomposition (shown in Figure A.3), which indicates a sequence of S and P that turns G into a single edge.*

We now proceed with the proof of Proposition 3. Using Proposition 7, we first verify whether G is series parallel, which can be done in linear time. If G is not series parallel, then it is not SLI as well. If G is series parallel, then

Figure A.2. (Color online) Two Operations That Turn a Series-Parallel Network into a Single Edge

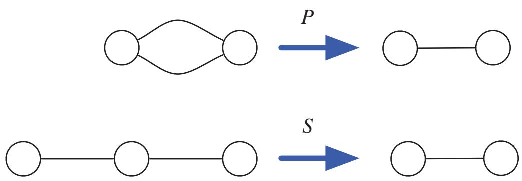

a binary tree decomposition can be obtained in linear time (again using Proposition 7). Note that the binary tree decomposition is not unique, and the following argument works with any binary tree decomposition. In this tree the edges of G are represented by the leaves of the tree. We label the incident edges to the origin by O and the incident edges to the destination by D (an edge might be labeled both O and D). Since G is SP, by definition, it is the result of attaching two SP networks in series or parallel. If it is the result of attaching two SP networks in a series, then there exists a node of the tree labeled S, referred to as the root of the tree, such that on one of the subtrees starting from that node we have only O labeled leaves and on the other subtree we have only D labeled leaves (this can be done in linear time by traversing the tree). If G is the result of attaching two SP networks in parallel, then there exists a node of the tree labeled P, again referred to as the root of the tree, such that on both subtrees starting from it we have both O and D labeled leaves.

We next show by induction on the size of tree that whether the binary tree represents an SLI network can be verified in linear time. If the root of the tree is S, then we have a series of two networks. By the induction hypothesis, in linear time we can verify whether each of these subtrees represents an SLI network, which in turn determines whether G is SLI. If the root of the tree is P, we need to check whether each subtree represents an LI network. We next show that this can be done in linear time, which concludes the proof.

Claim. *Given the binary tree decomposition, we can verify whether the underlying network is LI in linear time.*

We show this claim by induction on the size of the tree as well. Starting from the root of the tree, if the root has label P, then by the induction hypothesis, for each of the subtrees denoted by T_1 and T_2, we can verify whether the underlying network is LI in $O(V_{T_1})$ and $O(V_{T_2})$, respectively. The underlying network is LI if and only if both of these subtrees represent an LI network. Therefore, in $O(V)$ it can be verified whether the underlying network is LI. If the root is labeled S, then the underlying network is LI if and only if one of the subtrees is only labeled S and the other subtree is LI. Using any traversing algorithm (e.g., breadth-first search, depth-first search, etc.), one can visit all nodes in both subtrees in linear time, verifying if it only has S labels. Furthermore, by induction, we can verify whether each subtree represents an LI network. Therefore, in linear time, we can verify whether the network is LI, completing the proof.

A.3. Proofs of Section 5

A.3.1. Expansion of Example 2.

We provide an example for part (a) of Example 2. Let $K=1$, $c_{e_1}(x)=x$, $c_{e_2}(x)=1$, $c_{e_3}(x)=1$, $c_{e_4}(x)=x$, $c_{e_5}(x)=0$, and $s_1=1$. Also, we let the information sets be $\mathscr{E}_1=\{e_1,e_2,e_3,e_4\}$, and $\tilde{\mathscr{E}}_1=\{e_1,e_2,e_3,e_4,e_5\}$. In equilibrium, we have $f^{(1)}_{e_1e_3}=f^{(1)}_{e_2e_4}=\frac{1}{2}$ with $c^{(1)}=\frac{3}{2}$ and $\tilde{f}^{(1)}_{e_1e_3}=\tilde{f}^{(1)}_{e_2e_4}=0$, $\tilde{f}^{(1)}_{e_1e_5e_4}=1$ with $\tilde{c}^{(1)}=2$. Therefore, after expanding the information set of type 1 users, their equilibrium cost has increased from $\frac{3}{2}$ to 2.

A.3.2. Proof of the Claim of Remark 2.

We will show that there are infinitely many cost functions for the network shown in Figure 3(b) for which IBP occurs. In particular, we show the following claim.

Claim. *For any $a_1,a_3,a_5>0$ such that $a_1+a_3>a_5$, there exist nonnegative $b_1,b_2,b_3,b_4,b_5,a_2,s_1$, and s_2 such that with cost functions $c_{e_i}(x)=a_ix+b_i$, $1\leqslant i\leqslant 5$, IBP occurs in the network shown in Figure 3(b). In particular, we show that the following cost function parameters along with $\mathscr{E}_2=\{e_1,e_4,e_5\}$, $\mathscr{E}_1=\{e_2,e_3,e_5\}$, and $\tilde{\mathscr{E}}_1=\{e_1,e_2,e_3,e_5\}$ lead to IBP:*

$$a_4=b_1=b_3=b_5=0,$$
$$b_2=a_1y=a_1\frac{a_5(s_1+s_2)}{a_1+a_3+a_5},$$
$$b_4=\frac{a_5a_3(s_1+s_2)}{a_1+a_3+a_5},$$
$$\frac{s_1}{s_1+s_2}\in\left(\frac{a_1+a_3}{a_1+a_3+a_5},\ \min\left\{\frac{(a_3+a_5)(a_3a_5+a_1^2+a_1a_3+a_1a_5)-a_1a_5^2}{(a_1+a_3+a_5)(a_3a_5+a_3a_1+a_1a_5)},1\right\}\right),$$
$$a_2=\frac{a_5^2(s_1w-a_1)}{(a_5s_1-a_1s_2)w-a_3a_5}-a_3-a_5,\quad \textit{with}\quad w=\frac{a_1+a_3+a_5}{s_1+s_2}.$$

Proof. We let $a_4=b_1=b_3=b_5=0$ and then find a_2,b_2,b_4, s_1, and s_2 for which IBP occurs with $\mathscr{E}_2=\{e_1,e_4,e_5\}$, $\mathscr{E}_1=\{e_2,e_3,e_5\}$, and $\tilde{\mathscr{E}}_1=\{e_1,e_2,e_3,e_5\}$. We will find the a_2,b_2, b_4,s_1, and s_2 parameters such that before expanding the information set, the equilibrium flow is $f^{(2)}_{e_5}=0$, $f^{(2)}_{e_1e_4}=s_2$, and $f^{(1)}_{e_5}=s_1-x$, $f^{(1)}_{e_2e_3}=x$. We will further impose the constraint that the cost of route e_5 for type 2 users is equal to the cost of route e_1e_4. For this to hold, it is sufficient and necessary to have $a_5(s_1-x)=a_2x+b_2+a_3x$, which leads to

$$x=\frac{a_5s_1-b_2}{a_2+a_3+a_5}\in[0,s_1]. \tag{A.5}$$

We also have $a_1s_2+b_4=a_5(s_1-x)$, which leads to

$$a_1s_2+b_4=a_5\left(s_1-\frac{a_5s_1-b_2}{a_2+a_3+a_5}\right). \tag{A.6}$$

We will also choose a_2, b_2, b_4, s_1, and s_2 parameters such that after expanding the information set, the equilibrium flow becomes $\tilde{f}^{(2)}_{e_5}=s_2$, $\tilde{f}^{(2)}_{e_1e_4}=0$, $\tilde{f}^{(1)}_{e_5}=s_1-y$, $\tilde{f}^{(1)}_{e_2e_3}=0$, and $\tilde{f}^{(1)}_{e_1e_3}=y$. We will further impose the constraint that the cost of all available routes for each type of user is equal. For this to hold, it is sufficient and necessary to have $a_5(s_1+s_2-y)=a_1y+a_3y$, which leads to

$$y=\frac{a_5(s_2+s_1)}{a_1+a_3+a_5}\in(0,s_1). \tag{A.7}$$

We also have $a_1y+a_3y=b_2+a_3y$, which after substituting y from (A.7) leads to

$$b_2=a_1y=a_1\frac{a_5(s_2+s_1)}{a_1+a_3+a_5}. \tag{A.8}$$

Figure A.3. (Color online) Binary Tree Decomposition of a Series-Parallel Network

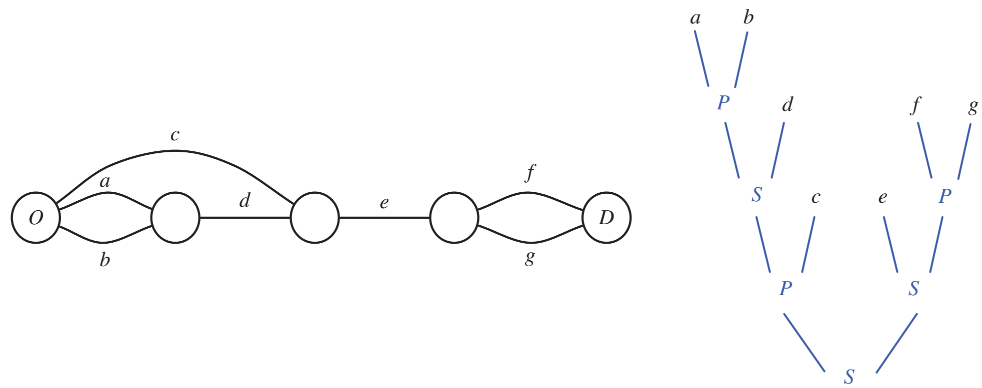

Also, for type 2 users we have $a_1 y + b_4 = a_5(s_2 + s_1 - y)$, which after substituting y from (A.7) leads to

$$b_4 = \frac{a_5 a_3 (s_2 + s_1)}{a_1 + a_3 + a_5}. \tag{A.9}$$

Therefore, Equations (A.9) and (A.8) determine b_2 and b_4 as a function of other parameters. In what follows we will show how to choose nonnegative s_1, s_2, and a_2 such that Equations (A.5)–(A.7) hold as well. After some rearrangements, we can see that the constraints imposed by Equations (A.5) and (A.7) are equivalent to

$$\frac{s_1}{s_2 + s_1} \geqslant \frac{\max\{a_5, a_1\}}{a_1 + a_3 + a_5}. \tag{A.10}$$

Furthermore, IBP occurs if we have $a_1 y + b_4 > a_1 s_2 + b_4$, which leads to $s_2/(s_2 + s_1) < a_5/(a_1 + a_3 + a_5)$ or, equivalently,

$$\frac{s_1}{s_2 + s_1} > \frac{a_1 + a_3}{a_1 + a_3 + a_5}. \tag{A.11}$$

Using $a_1 + a_3 > a_5$, Equations (A.10) and (A.11) become equivalent to

$$\frac{s_1}{s_2 + s_1} > \frac{a_1 + a_3}{a_1 + a_3 + a_5}. \tag{A.12}$$

Using Equation (A.6), we can find a_2 as follows:

$$a_2 = \frac{a_5^2 (s_1 w - a_1)}{a_5 s_1 w - a_5 a_3 + a_1 s_1 w - a_1(a_1 + a_3 + a_5)} - a_3 - a_5, \quad \text{where} \quad w = \frac{a_1 + a_3 + a_5}{s_1 + s_2}, \tag{A.13}$$

with the condition that the right-hand side of Equation (A.13) is nonnegative. From (A.12), the nonnegativity of a_2 becomes equivalent to

$$\frac{s_1}{s_1 + s_2}(a_1 + a_3 + a_5) \leqslant \frac{(a_3 + a_5)(a_3 a_5 + a_1^2 + a_1 a_3 + a_1 a_5) - a_1 a_5^2}{(a_3 + a_5)(a_5 + a_1) - a_5^2}. \tag{A.14}$$

Choosing $(s_1/(s_1 + s_2))(a_1 + a_3 + a_5)$, which satisfies both Equations (A.14) and (A.12), is feasible if we have

$$a_1 + a_3 < \frac{(a_3 + a_5)(a_3 a_5 + a_1^2 + a_1 a_3 + a_1 a_5) - a_1 a_5^2}{(a_3 + a_5)(a_5 + a_1) - a_5^2},$$

which after simplification becomes equivalent to $a_3 a_5^2 > 0$ and therefore holds. Hence, by choosing $(s_1/(s_1 + s_2))(a_1 + a_3 + a_5)$ such that

$$\frac{s_1}{s_1 + s_2}(a_1 + a_3 + a_5) \in \left(a_1 + a_3, \min\left\{\frac{(a_3 + a_5)(a_3 a_5 + a_1^2 + a_1 a_3 + a_1 a_5) - a_1 a_5^2}{(a_3 + a_5)(a_5 + a_1) - a_5^2}, a_1 + a_3 + a_5\right\}\right), \tag{A.15}$$

all the conditions are satisfied, and IBP occurs in this network for infinitely many cost functions. □

A.4. Proofs of Section 6

A.4.1. Proof of Lemma 1. Given the feasible flow $f^{(1:K)}$ for $(G, \mathscr{E}_{1:K}, s_{1:K}, \mathbf{c})$, we construct a feasible flow f with load $\sum_{i=1}^{K} s_i$ for a single type of user by letting $f_r = \sum_{i=1}^{K} f_r^{(i)}$. Using this construction, from two feasible flows $f^{(1:K)}$ and $\tilde{f}^{(1:K)}$, we obtain two feasible flows f and $\tilde{f}$ for a single-type congestion game such that the load of f is larger than or equal to the traffic demand of $\tilde{f}$. Therefore, part (a) follows from Milchtaich (2006, lemma 5).

We next show part (b). Since part (a) holds for any two feasible flows, we can apply it for the equilibrium flows $f^{(1:K)}$ and $\tilde{f}^{(1:K)}$ over the traffic networks $(G, \mathscr{E}_{1:K}, s_{1:K}, \mathbf{c})$ and $(G, \tilde{\mathscr{E}}_{1:K}, s_{1:K}, \mathbf{c})$, respectively (we can view $f^{(1:K)}$ as a feasible flow over the traffic network $(G, \tilde{\mathscr{E}}_{1:K}, s_{1:K}, \mathbf{c})$ as well). It follows that there exists a route r such that $\sum_{i=1}^{K} f_r^{(i)} > \sum_{i=1}^{K} \tilde{f}_r^{(i)}$ and $f_e \geqslant \tilde{f}_e$ for all $e \in r$. From the first inequality it follows that $\sum_{i=1}^{K} f_r^{(i)} > 0$, which shows at least one of the types (say, type i) sends positive traffic on route r. Note that i can be any element of $[K]$ (it can also be 1 as the flow $f^{(1:K)}$ is a feasible flow for the traffic network $(G, \mathcal{E}_{1:K}, s_{1:K}, \mathbf{c})$). We obtain

$$c^{(i)} = c_r \geqslant \tilde{c}_r \geqslant \tilde{c}^{(i)},$$

where the first equality follows from $f_r^{(i)} > 0$. The first inequality follows from $f_e \geqslant \tilde{f}_e$ for all $e \in r$. The second inequality follows from the definition of ICWE and the fact that if type i users can use route r in $(G, \mathscr{E}_{1:K}, s_{1:K}, \mathbf{c})$, then they can use it in $(G, \tilde{\mathscr{E}}_{1:K}, s_{1:K}, \mathbf{c})$ as well, since the information sets are not smaller in the second game. This completes the proof.

A.4.2. Proof of Lemma 2. *Part* (a): Suppose $f^{(1:K)}$ is an equilibrium flow on G. We next show that the restriction of $f^{(1:K)}$ to G_1 creates an equilibrium for G_1. Consider type i users and let r_1 be a route in G_1 such that $f_{r_1}^{(i)} > 0$, and let r_1' be another route in G_1, which belongs to the information set of type i users. The route r_1 is part of a route r in G for which $f_r^{(i)} > 0$. We let r_2 be the restriction of r to G_2 (so that $r = r_1 + r_2$). Since $f^{(1:K)}$ is an equilibrium of G, we have $c_r = c_{r_1} + c_{r_2} \leqslant c_{r_1'} + c_{r_2} = c_{r_1'+r_2}$, which leads to $c_{r_1} \leqslant c_{r_1'}$, showing that the restriction of $f^{(1:K)}$ to G_1 is an equilibrium. Similarly, the restriction to G_2 is an equilibrium.

Part (b): We consider an equilibrium $f^{(1:K)}$ for G, and then using part (a), we consider the equilibria of G_1 and G_2 obtained by restriction of $f^{(1:K)}$ to G_1 and G_2. For a type i and route r such that $f_r^{(i)} > 0$, we have $c^{(i)} = c_r = c_{r_1} + c_{r_2}$, where r_1 and r_2 are the restriction of r to G_1 and G_2, respectively (note that the only common node of r_1 and r_2 is the destination of G_1, which is the same as the origin of G_2; hence the operation $r_1 + r_2$ is a valid operation). Since $f_{r_1}^{(i)} > 0$ and $f_{r_2}^{(i)} > 0$, we have $c_{r_1} = c_1^{(i)}$ and $c_{r_2} = c_2^{(i)}$, which leads to $c^{(i)} = c_1^{(i)} + c_2^{(i)}$. □

A.4.3. Proof of Theorem 4. We first show two lemmas that we will use in the proof. The first lemma directly follows from the results of Milchtaich (2006) for a single-type congestion game.

Lemma 4. *Consider a traffic network with multiple information types $(G, \mathscr{E}_{1:K}, s_{1:K}, \mathbf{c})$. Let $f^{(1:K)}$ and $\tilde{f}^{(1:K)}$ be two (arbitrary) nonidentical feasible flows such that $\sum_{i=1}^{K} s_i \geqslant \sum_{i=1}^{K} \tilde{s}_i$. If G is series parallel, there exists a route r such that $f_e \geqslant \tilde{f}_e$ and $f_e > 0$ for all $e \in r$.*

Proof. Similar to the proof of Lemma 1, given a feasible flow $f^{(1:K)}$ for $(G, \mathscr{E}_{1:K}, s_{1:K}, \mathbf{c})$ we define a feasible flow f with traffic demand $\sum_{i=1}^{K} s_i$ for a congestion game with a single information type. Therefore, this lemma follows from Milchtaich (2006, lemma 2). □

Lemma 5. *Consider a traffic network with multiple information types $(G, \mathscr{E}_{1:K}, s_{1:K}, \mathbf{c})$, where $\mathscr{E}_i = \mathscr{E}$ for $i = 2, \dots, K$, $\mathscr{E}_1 \subseteq \mathscr{E}$, and G is a SP network. Consider an ICWE with flow $(f^{(1)}, \dots, f^{(K)})$ and let r be a route for which $f_e > 0$ for any $e \in r$. We have*

$$c_r \in \left[\min_{i\in[K]} c^{(i)}, \max_{i\in[K]} c^{(i)}\right],$$

where for any $i \in [K]$, $c^{(i)}$ denotes the equilibrium cost of type i users.

Proof. Since all the types except type 1 have full information, we have $c^{(i)} = c^{(j)}$ for all $i, j \in \{2, \dots, K\}$, $\max_{i\in[K]} c^{(i)} = c^{(1)}$, and $\min_{i\in[K]} c^{(i)} = c^{(j)}$, $j \neq 1$. By the definition of ICWE, we have $c_r \geqslant c^{(i)}$ (as $r \in \mathscr{R}_i$) for all $i \geqslant 2$. This leads to $c_r \geqslant \min_{i\in[K]} c^{(i)}$, showing the lower bound. We will next show the upper bound. We will prove this by induction on the number of edges of G. It evidently holds for a single edge as all equilibrium costs are equal to c_r. We next show the result for a series-parallel network G. Since G is SP, it is the result of either attaching two SP networks in series or attaching two SP networks in parallel. If G is the result of attaching two SP networks G_A and G_B in series, then using part (a) of Lemma 2, an ICWE for the overall network is obtained by concatenating an ICWE for G_A with an ICWE for G_B. We let r_A and r_B denote the sections of r that belong to G_A and G_B, respectively. We also let $c_A^{(i)}$ and $c_B^{(i)}$ be the equilibrium costs of type i users in G_A and G_B, respectively. By the induction hypothesis, we have $c_{r_A} \leqslant \max_{i\in[K]} c_A^{(i)}$ and $c_{r_B} \leqslant \max_{i\in[K]} c_B^{(i)}$. Since the traffic demands of type 1 users on both G_A and G_B are nonzero, we have $\max_{i\in[K]} c_A^{(i)} = c_A^{(1)}$ and $\max_{i\in[K]} c_B^{(i)} = c_B^{(1)}$. This leads to

$$c_r = c_{r_A} + c_{r_B} \leqslant c_A^{(1)} + c_B^{(1)} = c^{(1)},$$

where we used part (b) of Lemma 2 in the last equality.

Now suppose that G is the result of attaching G_A and G_B in parallel and suppose $r \in G_A$. Let $T = \{i \geqslant 2: f_A^{(i)} > 0\}$ denote the set of types that are sending a nonzero flow over G_A. Depending on whether $T = \emptyset$, we have the following two cases:

- $T = \emptyset$: Since $f_e > 0$ for all $e \in r$, at least one type must send a nonzero flow over G_A and since $T = \emptyset$, only type 1 sends a nonzero flow over G_A. Therefore, we have $c_r = c^{(1)}$. We also have $c^{(1)} \leqslant \max_{i\in[K]} c^{(i)}$, leading to $c_r \leqslant \max_{i\in[K]} c^{(i)}$.
- $T \neq \emptyset$: We either have $f_A^{(1)} > 0$ or $f_A^{(1)} = 0$. If $f_A^{(1)} > 0$, then by the induction hypothesis, we have

$$c_r \leqslant \max_{i\in T\cup\{1\}} c_A^{(i)} = \max_{i\in T\cup\{1\}} c^{(i)} \leqslant \max_{i\in[K]} c^{(i)},$$

where the equality holds because each type $i \in T \cup \{1\}$ sends a positive flow over A and its equilibrium cost in G is the same as its equilibrium cost in G_A. If $f_A^{(1)} = 0$, then again by the induction hypothesis and using $T \neq \emptyset$, we have

$$c_r \leqslant \max_{i\in T} c_A^{(i)} = \max_{i\in T} c^{(i)} \leqslant \max_{i\in[K]} c^{(i)},$$

where the equality holds because each type $i \in T$ sends a positive flow over A.

This concludes the proof of lemma.

Proof of Part (a) of Theorem 4. After expanding information set of type 1 users to $\mathscr{E}$, we obtain $\tilde{c}^{(i)} = \tilde{c}^{(1)}$ for all $i \in [K]$. Using Lemma 4, there exists a route r such that $f_e \geqslant \tilde{f}_e$ and $f_e > 0$ for any $e \in r$. We have

$$c_r \geqslant \tilde{c}_r \geqslant \tilde{c}^{(i)} = \tilde{c}^{(1)}, \quad \forall\, i \in [K],$$

where the first inequality follows from $f_e \geqslant \tilde{f}_e$, the second inequality follows from the definition of ICWE, and the equality follows from $\tilde{\mathscr{E}}_i = \mathscr{E}$ for all $i = 1, \dots, K$. Since $\mathscr{E}_1 \subseteq \mathscr{E}$, we have $c^{(i)} = c^{(j)} \leqslant c^{(1)}$ for all $i, j = 2, \dots, K$. Using Lemma 5, this leads to

$$c_r \leqslant \max_{i\in[K]} c^{(i)} = c^{(1)}.$$

Combining the previous two relations leads to $\tilde{c}^{(1)} \leqslant c^{(1)}$.

Proof of Part (b) of Theorem 4. The proof is similar to the proof of part (b) of Theorem 3. In Example 2, we have provided an example showing that IBP with restricted information sets can occur over the Wheatstone network shown in Figure 3(a).

Suppose that a network G is not series parallel. Using Proposition 2, G can be constructed from the Wheatstone network shown in Figure 3(a) by following the steps of embedding. To construct an example for G, we start from the cost functions for which the embedded network features IBP with restricted information sets, and then following the steps of embedding, we will update the information sets as well as the cost functions in a way that IBP occurs in the final network, which is G. The updates are identical to those described in the proof of part (b) of Theorem 3 and establishes that if IBP with restricted information sets is present in the initial network (i.e., the Whetstone network shown in Figure 3(a)), it will be present in network G as well. This completes the proof of part (b).

A.4.4. Omitted Proof of Example 3. First, note that after expansion of information, without loss of generality, each type of user (i,j), $i=1,2$ has at least two routes from O_i to D_i. Because, otherwise, if a type with traffic demand s has only one route r, we can consider an equivalent game in which we update the cost of all edges on r from $c_e(x)$ to $c_e(x+s)$. Also, note that because of symmetry, we can only consider the information expansion of one of the types of the form $(1,j)$. Therefore, without loss of generality, we assume that there exists one type from O_2 to D_2 with information about all edges of the network and there exists either one or two types from O_1 to D_1. Below, we examine all possible cases and show that IBP does not occur:

1. There exist two types $\{(1,1),(2,1)\}$ such that $\mathcal{R}_{2,1}=\{e_1e_3,e_2\}$, $\mathcal{R}_{1,1}=\{e_1\}$, and $\tilde{\mathcal{R}}_{1,1}=\{e_1,e_2e_3\}$: If type $(1,1)$ does not use route e_2e_3 after information expansion, then the equilibrium remains the same. Now suppose that type $(1,1)$ uses route e_2e_3 (i.e., $\tilde{f}^{(1,1)}_{e_1}<f^{(1,1)}_{e_1}=s_{1,1}$). If $\tilde{f}_{e_1}\leqslant f_{e_1}$, then we have

$$\tilde{c}^{(1,1)}\leqslant c_{e_1}(\tilde{f}_{e_1})\leqslant c_{e_1}(f_{e_1})=c^{(1,1)},$$

which shows that IBP does not occur. Now suppose $\tilde{f}_{e_1}>f_{e_1}$, which in turn shows $\tilde{f}_{e_2}<f_{e_2}$ as $\tilde{f}_{e_1}+\tilde{f}_{e_2}=f_{e_1}+f_{e_2}=s_{1,1}+s_{2,1}$.
We have

$$\tilde{f}_{e_3}=\tilde{f}^{(1,1)}_{e_2e_3}+\tilde{f}^{(2,1)}_{e_1e_3}>f^{(1,1)}_{e_2e_3}+f^{(2,1)}_{e_1e_3}=f_{e_3},$$

where we used $\tilde{f}^{(1,1)}_{e_2e_3}>f^{(1,1)}_{e_2e_3}=0$ and $\tilde{f}^{(2,1)}_{e_1e_3}\geqslant f^{(2,1)}_{e_1e_3}$, which holds because $\tilde{f}^{(2,1)}_{e_1e_3}=\tilde{f}_{e_1}-\tilde{f}^{(1,1)}_{e_1}>f_{e_1}-s_{1,1}=f_{e_1}-f^{(1,1)}_{e_1}=f^{(2,1)}_{e_1e_3}$. Therefore, we have

$$c_{e_1}(\tilde{f}_{e_1})+c_{e_3}(\tilde{f}_{e_3})\leqslant c_{e_2}(\tilde{f}_{e_2})\leqslant c_{e_2}(f_{e_2})\leqslant c_{e_1}(f_{e_1})+c_{e_3}(f_{e_3}),\quad (A.16)$$

where the first inequality holds because $\tilde{f}^{(2,1)}_{e_1e_3}>f^{(2,1)}_{e_1e_3}\geqslant 0$, the second inequality holds because $\tilde{f}_{e_2}<f_{e_2}$, and the third inequality holds because $f^{(2,1)}_{e_2}=s_{2,1}-f^{(2,1)}_{e_1e_3}>s_{2,1}-\tilde{f}^{(2,1)}_{e_1e_3}=\tilde{f}^{(2,1)}_{e_2}\geqslant 0$. Inequality (A.16), together with $c_{e_1}(\tilde{f}_{e_1})\geqslant c_{e_1}(f_{e_1})$ and $c_{e_3}(\tilde{f}_{e_3})\geqslant c_{e_3}(f_{e_3})$, shows that the cost of all three edges before and after information expansion are the same, leading to the same equilibrium cost for all types. Therefore, IBP does not occur in this case.

2. There exist two types $\{(1,1),(2,1)\}$ such that $\mathcal{R}_{2,1}=\{e_1e_3,e_2\}$, $\mathcal{R}_{1,1}=\{e_2e_3\}$, and $\tilde{\mathcal{R}}_{1,1}=\{e_1,e_2e_3\}$: If type $(1,1)$ does not use e_1 after information expansion, then the equilibrium remains the same. Now suppose type $(1,1)$ uses route e_1. We show that IBP does not occur in this case by considering all possibilities as follows:

a. $\tilde{f}_{e_1}\leqslant f_{e_1}$: Since $f_{e_1}+f_{e_2}=s_{1,1}+s_{2,1}=\tilde{f}_{e_1}+\tilde{f}_{e_2}$, we have $\tilde{f}_{e_2}\geqslant f_{e_2}$. We also have $\tilde{f}_{e_3}<f_{e_3}$, because

$$\tilde{f}_{e_3}=\tilde{f}^{(1,1)}_{e_2e_3}+\tilde{f}^{(2,1)}_{e_1e_3}<f^{(1,1)}_{e_2e_3}+f^{(2,1)}_{e_1e_3}=f_{e_3},$$

where we used $\tilde{f}^{(1,1)}_{e_2e_3}<f^{(1,1)}_{e_2e_3}$ as type $(1,1)$ is using e_1 after information expansion and $\tilde{f}^{(2,1)}_{e_1e_3}<f^{(2,1)}_{e_1e_3}$ as $\tilde{f}^{(2,1)}_{e_1e_3}=\tilde{f}_{e_1}-\tilde{f}^{(1,1)}_{e_1}<f_{e_1}=f^{(2,1)}_{e_1e_3}$. The inequality $\tilde{f}^{(2,1)}_{e_1e_3}<f^{(2,1)}_{e_1e_3}$ implies $\tilde{f}^{(2,1)}_{e_2}>f^{(2,1)}_{e_2}\geqslant 0$. Therefore, we have

$$c_{e_2}(\tilde{f}_{e_2})\leqslant c_{e_1}(\tilde{f}_{e_1})+c_{e_3}(\tilde{f}_{e_3})\leqslant c_{e_1}(f_{e_1})+c_{e_3}(f_{e_3})\leqslant c_{e_2}(f_{e_2}),\quad (A.17)$$

where the first inequality follows from $\tilde{f}^{(2,1)}_{e_2}>0$, the second inequality follows from $\tilde{f}_{e_1}\leqslant f_{e_1}$ and $\tilde{f}_{e_3}\leqslant f_{e_3}$, and the third inequality follows from $f^{(2,1)}_{e_1e_3}>\tilde{f}^{(2,1)}_{e_1e_3}\geqslant 0$. Inequality (A.17) leads to

$$c^{(1,1)}=c_{e_2}(\tilde{f}_{e_2})+c_{e_3}(\tilde{f}_{e_3})\leqslant c_{e_2}(f_{e_2})+c_{e_3}(f_{e_3})\leqslant \tilde{c}^{(1,1)},$$

showing that IBP does not occur.

b. $\tilde{f}_{e_1}>f_{e_1}$: Since $f_{e_1}+f_{e_2}=s_{1,1}+s_{2,1}=\tilde{f}_{e_1}+\tilde{f}_{e_2}$, we have $\tilde{f}_{e_2}<f_{e_2}$. If $\tilde{f}_{e_3}\leqslant f_{e_3}$, then we have

$$c^{(1,1)}=c_{e_2}(\tilde{f}_{e_2})+c_{e_3}(\tilde{f}_{e_3})\leqslant c_{e_2}(f_{e_2})+c_{e_3}(f_{e_3})\leqslant \tilde{c}^{(1,1)},$$

showing that IBP does not occur. Otherwise, we have $\tilde{f}_{e_3}>f_{e_3}$. First note that if $\tilde{f}_{e_3}>f_{e_3}$, then $\tilde{f}^{(2,1)}_{e_1e_3}>f^{(2,1)}_{e_1e_3}$. This inequality holds because

$$\tilde{f}^{(2,1)}_{e_1e_3}=\tilde{f}_{e_3}-\tilde{f}^{(1,1)}_{e_2e_3}>f_{e_3}-f^{(1,1)}_{e_2e_3}=f^{(2,1)}_{e_1e_3},$$

where we used $\tilde{f}_{e_3}>f_{e_3}$ and $\tilde{f}^{(1,1)}_{e_2e_3}<f^{(1,1)}_{e_2e_3}$ as $(1,1)$ uses e_1 after the expansion of information (i.e., $f^{(1,1)}_{e_2e_3}=s_{1,1}$ and $\tilde{f}^{(1,1)}_{e_2e_3}<s_{1,1}$). Therefore, we have

$$c_{e_1}(\tilde{f}_{e_1})+c_{e_3}(\tilde{f}_{e_3})\leqslant c_{e_2}(\tilde{f}_{e_2})\leqslant c_{e_2}(f_{e_2})\leqslant c_{e_1}(f_{e_1})+c_{e_3}(f_{e_3}),\quad (A.18)$$

where the first inequality holds because $\tilde{f}^{(2,1)}_{e_1e_3}>f^{(2,1)}_{e_1e_3}\geqslant 0$, the second inequality holds because $\tilde{f}_{e_2}<f_{e_2}$, and the third inequality holds because $f^{(2,1)}_{e_2}=s_{2,1}-f^{(2,1)}_{e_1e_3}>s_{2,1}-\tilde{f}^{(2,1)}_{e_1e_3}=\tilde{f}^{(2,1)}_{e_2}\geqslant 0$. Inequality (A.18), together with $\tilde{f}_{e_1}>f_{e_1}$ and $\tilde{f}_{e_3}>f_{e_3}$, leads to $c_{e_1}(f_{e_1})=c_{e_1}(\tilde{f}_{e_1})$, $c_{e_2}(f_{e_2})=c_{e_2}(\tilde{f}_{e_2})$, and $c_{e_3}(f_{e_3})=c_{e_3}(\tilde{f}_{e_3})$. Therefore, we have

$$\tilde{c}^{(1,1)}\leqslant c_{e_2}(\tilde{f}_{e_2})+c_{e_3}(\tilde{f}_{e_3})=c_{e_2}(f_{e_2})+c_{e_3}(f_{e_3})=c^{(1,1)},$$

showing that IBP does not occur.

3. There exist three types $\{(1,1),(1,2),(2,1)\}$ such that $\mathcal{R}_{2,1}=\{e_1e_3,e_2\}$, $\mathcal{R}_{1,2}=\{e_1,e_2e_3\}$, $\mathcal{R}_{1,1}=\{e_1\}$, and $\tilde{\mathcal{R}}_{1,1}=\{e_1,e_2e_3\}$: This case is similar to the first case. If type $(1,1)$ does not use e_2e_3 after the expansion of information, then the equilibrium remains the same. Now suppose that type $(1,1)$ uses route e_2e_3 (i.e., $\tilde{f}^{(1,1)}_{e_1}<f^{(1,1)}_{e_1}=s_{1,1}$). If $\tilde{f}_{e_1}\leqslant f_{e_1}$, then we have

$$\tilde{c}^{(1,1)}\leqslant c_{e_1}(\tilde{f}_{e_1})\leqslant c_{e_1}(f_{e_1})=c^{(1,1)},$$

which shows that IBP does not occur. Now suppose that $\tilde{f}_{e_1}>f_{e_1}$, which in turn shows $\tilde{f}_{e_2}<f_{e_2}$ as $\tilde{f}_{e_1}+\tilde{f}_{e_2}=f_{e_1}+f_{e_2}=s_{1,1}+s_{1,2}+s_{2,1}$. We consider the following two cases:

a. $\tilde{f}^{(1,1)}_{e_1}+\tilde{f}^{(1,2)}_{e_1}<f^{(1,1)}_{e_1}+f^{(1,2)}_{e_1}$ (note that $f^{(1,1)}_{e_1}=s_{1,1}$): We have

$$\tilde{f}^{(2,1)}_{e_1e_3}=\tilde{f}_{e_1}-(\tilde{f}^{(1,1)}_{e_1}+\tilde{f}^{(1,2)}_{e_1})>f_{e_1}-(f^{(1,1)}_{e_1}+f^{(1,2)}_{e_1})=f^{(2,1)}_{e_1e_3},$$
$$\tilde{f}^{(1,1)}_{e_2e_3}+\tilde{f}^{(1,2)}_{e_2e_3}=s_{1,1}+s_{1,2}-(\tilde{f}^{(1,1)}_{e_1}+\tilde{f}^{(1,2)}_{e_1})>s_{1,1}+s_{1,2}-(f^{(1,1)}_{e_1}+f^{(1,2)}_{e_1})=f^{(1,2)}_{e_2e_3}.$$

These two inequalities lead to

$$\tilde{f}_{e_3}=\tilde{f}^{(2,1)}_{e_1e_3}+\tilde{f}^{(1,1)}_{e_2e_3}+\tilde{f}^{(1,2)}_{e_2e_3}>f^{(2,1)}_{e_1e_3}+f^{(1,2)}_{e_2e_3}=f_{e_3}.$$

Therefore, we have

$$c_{e_1}(\tilde{f}_{e_1})+c_{e_3}(\tilde{f}_{e_3})\leqslant c_{e_2}(\tilde{f}_{e_2})\leqslant c_{e_2}(f_{e_2})\leqslant c_{e_1}(f_{e_1})+c_{e_3}(f_{e_3}), \quad \text{(A.19)}$$

where the first inequality holds because $\tilde{f}^{(2,1)}_{e_1e_3} > f^{(2,1)}_{e_1e_3} \geqslant 0$, the second inequality holds because $\tilde{f}_{e_2} < f_{e_2}$, and the third inequality holds because $f^{(2,1)}_{e_2} = s_{2,1} - f^{(2,1)}_{e_1e_3} > s_{2,1} - \tilde{f}^{(2,1)}_{e_1e_3} = \tilde{f}^{(2,1)}_{e_2} \geqslant 0$. Inequality (A.19), together with $\tilde{f}_{e_1} > f_{e_1}$ and $\tilde{f}_{e_3} > f_{e_3}$, leads to $c_{e_1}(f_{e_1}) = c_{e_1}(\tilde{f}_{e_1})$, $c_{e_2}(f_{e_2}) = c_{e_2}(\tilde{f}_{e_2})$, and $c_{e_3}(f_{e_3}) = c_{e_3}(\tilde{f}_{e_3})$. Therefore, the cost of all three edges before and after information expansion are the same, leading to the same equilibrium cost for all types. Therefore, IBP does not occur in this case.

b. $\tilde{f}^{(1,1)}_{e_1} + \tilde{f}^{(1,2)}_{e_1} \geqslant f^{(1,1)}_{e_1} + f^{(1,2)}_{e_1}$: We have

$$\begin{aligned}\tilde{f}^{(1,2)}_{e_1} &= (\tilde{f}^{(1,1)}_{e_1} + \tilde{f}^{(1,2)}_{e_1}) - \tilde{f}^{(1,1)}_{e_1} \geqslant (f^{(1,1)}_{e_1} + f^{(1,2)}_{e_1}) - \tilde{f}^{(1,1)}_{e_1}\\ &> (f^{(1,1)}_{e_1} + f^{(1,2)}_{e_1}) - f^{(1,1)}_{e_1} = f^{(1,2)}_{e_1}, \qquad \text{(A.20)}\end{aligned}$$

$$\begin{aligned}\tilde{f}^{(1,1)}_{e_2e_3} + \tilde{f}^{(1,2)}_{e_2e_3} &= s_{1,1} + s_{1,2} - (\tilde{f}^{(1,1)}_{e_1} + \tilde{f}^{(1,2)}_{e_1}) \leqslant s_{1,1} + s_{1,2}\\ &\quad - (f^{(1,1)}_{e_1} + f^{(1,2)}_{e_1}) = f^{(1,2)}_{e_2e_3}. \qquad \text{(A.21)}\end{aligned}$$

If $\tilde{f}^{(2,1)}_{e_1e_3} \leqslant f^{(2,1)}_{e_1e_3}$, then inequality (A.21) leads to

$$\tilde{f}_{e_3} = \tilde{f}^{(1,1)}_{e_2e_3} + \tilde{f}^{(1,2)}_{e_2e_3} + \tilde{f}^{(2,1)}_{e_1e_3} \leqslant f^{(1,2)}_{e_2e_3} + f^{(2,1)}_{e_1e_3} = f_{e_3}.$$

Therefore, we obtain

$$c_{e_1}(\tilde{f}_{e_1}) = c_{e_2}(\tilde{f}_{e_2}) + c_{e_3}(\tilde{f}_{e_3}) \leqslant c_{e_2}(f_{e_2}) + c_{e_3}(f_{e_3}) \leqslant c_{e_1}(f_{e_1}),$$

where the first equality holds because, using inequality (A.20), we obtain $\tilde{f}^{(1,2)}_{e_1} > f^{(1,2)}_{e_1} \geqslant 0$ and $\tilde{f}^{(1,1)}_{e_2e_3} > 0$, the first inequality holds because $\tilde{f}_{e_2} < f_{e_2}$ and $\tilde{f}_{e_3} < f_{e_3}$, and the second inequality holds because, using inequality (A.20), we obtain $f^{(1,2)}_{e_2e_3} = s_{1,2} - f^{(1,2)}_{e_1} > s_{1,2} - \tilde{f}^{(1,2)}_{e_1} \geqslant 0$. Inequality (A.22) leads to

$$\tilde{c}^{(1,1)} \leqslant c_{e_1}(\tilde{f}_{e_1}) \leqslant c_{e_1}(f_{e_1}) = c^{(1,1)},$$

showing that IBP does not occur in this case.

Now suppose that $\tilde{f}^{(2,1)}_{e_1e_3} > f^{(2,1)}_{e_1e_3}$, which leads to

$$\begin{aligned}c_{e_1}(\tilde{f}_{e_1}) &\leqslant c_{e_1}(\tilde{f}_{e_1}) + c_{e_3}(\tilde{f}_{e_3}) \leqslant c_{e_2}(\tilde{f}_{e_2}) \leqslant c_{e_2}(f_{e_2})\\ &\leqslant c_{e_2}(f_{e_2}) + c_{e_3}(f_{e_3}) \leqslant c_{e_1}(f_{e_1}), \qquad \text{(A.22)}\end{aligned}$$

where the second inequality holds because $\tilde{f}^{(2,1)}_{e_1e_3} > f^{(2,1)}_{e_1e_3} \geqslant 0$, the third inequality holds because $\tilde{f}_{e_2} < f_{e_2}$, and the last inequality holds because, using inequality (A.20), we obtain $f^{(1,2)}_{e_2e_3} = s_{1,2} - f^{(1,2)}_{e_1} > s_{1,2} - \tilde{f}^{(1,2)}_{e_1} \geqslant 0$. Inequality (A.22) leads to

$$\tilde{c}^{(1,1)} \leqslant c_{e_1}(\tilde{f}_{e_1}) \leqslant c_{e_1}(f_{e_1}) = c^{(1,1)},$$

showing that IBP does not occur in this case.

(4) There exist three types $\{(1,1),(1,2),(2,1)\}$ such that $\mathcal{R}_{2,1} = \{e_1e_3, e_2\}$, $\mathcal{R}_{1,2} = \{e_1, e_2e_3\}$, $\mathcal{R}_{1,1} = \{e_2e_3\}$, and $\tilde{\mathcal{R}}_{1,1} = \{e_1, e_2e_3\}$: This case is similar to the second case.

References

Acemoglu D, Ozdaglar A (2007) Competition and efficiency in congested markets. *Math. Oper. Res.* 32(1):1–31.

Acemoglu D, Johari R, Ozdaglar A (2007) Partially optimal routing. *IEEE J. Selected Areas Comm.* 25(6):1148–1160.

Anshelevich E, Dasgupta A, Kleinberg J, Tardos É, Wexler T, Roughgarden T (2008) The price of stability for network design with fair cost allocation. *SIAM J. Comput.* 38(4):1602–1623.

Arnott R, Small K (1994) The economics of traffic congestion. *Amer. Scientist* 82(5):446–455.

Arnott R, De Palma A, Lindsey R (1991) Does providing information to drivers reduce traffic congestion? *Transportation Res. Part A: General* 25(5):309–318.

Ashlagi I, Monderer D, Tennenholtz M (2009) Two-terminal routing games with unknown active players. *Artificial Intelligence* 173(15):1441–1455.

Beckmann M, McGuire CB, Winsten CB (1956) Studies in the economics of transportation. Research Memorandum RM-1488, RAND Corporation, Santa Monica, CA.

Ben-Akiva M, De Palma A, Isam K (1991) Dynamic network models and driver information systems. *Transportation Res. Part A: General* 25(5):251–266.

Bertsekas DP (1999) *Nonlinear Programming* (Athena Scientific, Nashua, NH).

Bondy JA, Murty USR (1976) *Graph Theory with Applications*, Vol. 290 (Macmillan, London).

Bottom J, Ben-Akiva M, Bierlaire M, Chabini I, Koutsopoulos H, Yang Q (1999) Investigation of route guidance generation issues by simulation with dynaMIT. *Proc. 14th Internat. Sympos. Transportation Traffic Theory* (Pergamon, Oxford, UK), 577–600.

Braess D (1968) Über ein Paradoxon aus der Verkehrsplanung. *Math. Methods Oper. Res.* 12(1):258–268.

Chen X, Diao Z, Hu X (2015) Excluding Braess's paradox in nonatomic selfish routing. Hoefer M, ed. *Algorithmic Game Theory*, Lecture Notes in Computer Science, Vol. 9347 (Springer, Berlin), 219–230.

Chen X, Diao Z, Hu X (2016) Network characterizations for excluding Braess paradox. *Theory Comput. Systems* 59(4):747–780.

Cohen JE, Horowitz P (1991) Paradoxical behaviour of mechanical and electrical networks. *Nature* 352:699–701.

Cohen JE, Jeffries C (1997) Congestion resulting from increased capacity in single-server queueing networks. *IEEE/ACM Trans. Networking* 5(2):305–310.

Correa JR, Schulz AS, Stier-Moses NE (2004) Selfish routing in capacitated networks. *Math. Oper. Res.* 29(4):961–976.

Correa JR, Schulz AS, Stier-Moses NE (2005) On the inefficiency of equilibria in congestion games. Cook WJ, Schulz AS, eds. *Integer Programming and Combinatorial Optimization*, Lecture Notes in Computer Science, Vol. 2337 (Springer, Berlin), 167–181.

Dafermos S, Nagurney A (1984) On some traffic equilibrium theory paradoxes. *Transportation Res. Part B: Methodological* 18(2):101–110.

Diestel R (2000) *Graph Theory*, Graduate Texts in Mathematics, Vol. 173 (Springer, New York).

Duffin RJ (1965) Topology of series-parallel networks. *J. Math. Anal. Appl.* 10(2):303–318.

Dughmi S (2014) On the hardness of signaling. *Proc. 2014 IEEE 55th Annual Sympos. Foundations Comput. Sci.* (*FOCS*) (IEEE Computer Society, Washington, DC), 354–363.

Eppstein D (1992) Parallel recognition of series-parallel graphs. *Inform. Comput.* 98(1):41–55.

Epstein A, Feldman M, Mansour Y (2009) Efficient graph topologies in network routing games. *Games Econom. Behav.* 66(1):115–125.

Feldman M, Friedler O (2015) Convergence to strong equilibrium in network design games. *ACM SIGMETRICS Performance Evaluation Rev.* 43(3):71–71.

Frank M (1981) The Braess paradox. *Math. Programming* 20(1):283–302.

Friedman EJ (2004) Genericity and congestion control in selfish routing. *Proc. 43rd IEEE Conf. Decision Control*, Vol. 5 (IEEE Computer Society, Los Alamitos, CA), 4667–4672.

Fujishige S, Goemans MX, Harks T, Peis B, Zenklusen R (2017) Matroids are immune to Braess' paradox. *Math. Oper. Res.* 42(3):745–761.

Gairing M, Klimm M (2013) Congestion games with player-specific costs revisited. Vöcking B, ed. *Internat. Sympos. Algorithmic Game Theory* (Springer, Berlin), 98–109.

Gairing M, Monien B, Tiemann K (2006) Routing (un-)splittable flow in games with player-specific linear latency functions. Bugliesi M, Preneel B, Sassone V, Wegener I, eds. *Automata, Languages, Programming 33rd Internat. Colloquium* (Springer, Berlin), 501–512.

Georgiou C, Pavlides T, Philippou A (2009) Selfish routing in the presence of network uncertainty. *Parallel Processing Lett.* 19(1):141–157.

Harel M, Mossel E, Strack P, Tamuz O (2014) Groupthink and the failure of information aggregation in large groups. Working paper, University of Geneva, Geneva. http://dx.doi.org/10.2139/ssrn.2541707.

Holzman R, Law-yone N (2003) Network structure and strong equilibrium in route selection games. *Math. Soc. Sci.* 46(2):193–205.

Holzman R, Monderer D (2015) Strong equilibrium in network congestion games: Increasing versus decreasing costs. *Internat. J. Game Theory* 44(3):647–666.

Jahn O, Möhring RH, Schulz AS, Stier-Moses NE (2005) System-optimal routing of traffic flows with user constraints in networks with congestion. *Oper. Res.* 53(4):600–616.

Johari R, Tsitsiklis J (2004) Network resource allocation and a congestion game. *Math. Oper. Res.* 29(3):407–435.

Kelly FP, Maulloo AK, Tan DKH (1998) Rate control for communication networks: Shadow prices, proportional fairness and stability. *J. Oper. Res. Soc.* 49(3):237–252.

Konishi H, Le Breton M, Weber S (1997) Equilibria in a model with partial rivalry. *J. Econom. Theory* 72(1):225–237.

Korilis YA, Lazar AA, Orda A (1997) Achieving network optima using Stackelberg routing strategies. *IEEE/ACM Trans. Networking* 5(1):161–173.

Koutsoupias E, Papadimitriou C (1999) Worst-case equilibria. Meinel C, Tilson S, eds. *Proc. 16th Annual Conf. Theoret. Aspects Comput. Sci.* (*STACS 99*) (Springer, Berlin), 404–413.

Lin H, Roughgarden T, Tardos É (2004) A stronger bound on Braess's paradox. *Proc. 15th Annual ACM-SIAM Sympos. Discrete Algorithms* (Society for Industrial and Applied Mathematics, Philadelphia), 340–341.

Lin H, Roughgarden T, Tardos É, Walkover A (2011) Stronger bounds on Braess's paradox and the maximum latency of selfish routing. *SIAM J. Discrete Math.* 25(4):1667–1686.

Liu J, Amin S, Schwartz G (2016) Effects of information heterogeneity in Bayesian routing games. Working paper, Massachusetts Institute of Technology, Cambridge. arXiv:1603.08853.

Low SH, Lapsley DE (1999) Optimization flow control: Basic algorithm and convergence. *IEEE/ACM Trans. Networking* 7(6): 861–874.

Maheswaran RT, Başar T (2003) Nash equilibrium and decentralized negotiation in auctioning divisible resources. *Group Decision Negotiation* 12(5):361–395.

Mahmassani H, Herman R (1984) Dynamic user equilibrium departure time and route choice on idealized traffic arterials. *Transportation Sci.* 18(4):362–384.

Mavronicolas M, Milchtaich I, Monien B, Tiemann K (2007) Congestion games with player-specific constants. Kučera L, Kučera A, eds. *Mathematical Foundations of Computer Science* 2007, Lecture Notes in Computer Science, Vol. 4708 (Springer, Berlin), 633–644.

Meir R, Parkes D (2014) Playing the wrong game: Smoothness bounds for congestion games with risk averse agents. Working paper, Technion–Isreal Institute of Technology, Haifa. arXiv:1411.1751.

Meir R, Parkes D (2015) Playing the wrong game: Smoothness bounds for congestion games with behavioral biases. *ACM SIGMETRICS Performance Evaluation Rev.* 43(3):67–70.

Milchtaich I (2000) Generic uniqueness of equilibrium in large crowding games. *Math. Oper. Res.* 25(3):349–364.

Milchtaich I (2004a) Random-player games. *Games Econom. Behav.* 47(2):353–388.

Milchtaich (2004b) Social optimality and cooperation in nonatomic congestion games. *J. Econom. Theory* 114(1):56–87.

Milchtaich I (2005) Topological conditions for uniqueness of equilibrium in networks. *Math. Oper. Res.* 30(1):225–244.

Milchtaich I (2006) Network topology and the efficiency of equilibrium. *Games Econom. Behav.* 57(2):321–346.

Murchland JD (1970) Braess's paradox of traffic flow. *Transportation Res.* 4(4):391–394.

Newell GF (1980) *Traffic Flow on Transportation Networks*, MIT Press Series in Transportation Studies, Vol. 5 (MIT Press, Cambridge, MA).

Nikolova E, Stier-Moses N (2015) The burden of risk aversion in mean-risk selfish routing. *Proc. of the 16th ACM Conference on Economics and Computation* (ACM, New York, NY), 489–506.

Nisan N, Roughgarden T, Tardos É, Vazirani VV (2007) *Algorithmic Game Theory*, Vol. 1 (Cambridge University Press, Cambridge, UK).

Orda A, Rom R, Shimkin N (1993) Competitive routing in multiuser communication networks. *IEEE/ACM Trans. Networking* 1(5):510–521.

Ordóñez F, Stier-Moses NE (2010) Wardrop equilibria with risk-averse users. *Transportation Sci.* 44(1):63–86.

Patriksson M (1994) *The Traffic Assignment Problem: Models and Methods* (VSP Publishers, Utrecht, Netherlands).

Perakis G (2004) The price of anarchy when costs are non-separable and asymmetric. Bienstock D, Nemhauser G, eds. *Integer Programming and Combinatorial Optimization* (Springer, Berlin), 46–58.

Pigou AC (1920) *The Economics of Welfare* (Macmillan and Co., London).

Riordan J, Shannon C (1942) The number of two-terminal series-parallel networks. *J. Math. Phys.* 21(1–4):83–93.

Rogers R, Roth A, Ullman J, Wu ZS (2015) Inducing approximately optimal flow using truthful mediators. *Proc. of the 16th ACM Conference on Economics and Computation* (ACM, New York, NY), 471–488.

Rosenthal RW (1973) A class of games possessing pure-strategy Nash equilibria. *Internat. J. Game Theory* 2(1):65–67.

Roughgarden T (2001) Designing networks for selfish users is hard. *Foundations Comput. Sci. 2001: Proc. 42nd IEEE Sympos.* (IEEE, Los Alamitos, CA), 472–481.

Roughgarden T (2006) On the severity of Braess's paradox: Designing networks for selfish users is hard. *J. Comput. System Sci.* 72(5):922–953.

Roughgarden T, Tardos É (2002) How bad is selfish routing? *J. ACM* 49(2):236–259.

Samuelson PA (1952) Spatial price equilibrium and linear programming. *Amer. Econom. Rev.* 42(3):283–303.

Sanghavi S, Hajek B (2004) Optimal allocation of a divisible good to strategic buyers. *Proc. 43rd IEEE Conf. Decision Control*(*CDC*), Vol. 3 (IEEE Computer Society, Los Alamitos, CA), 2748–2753.

Schmeidler D (1973) Equilibrium points of nonatomic games. *J. Statist. Phys.* 7(4):295–300.

Schoenmakers LAM (1995) *A New Algorithm for the Recognition of Series Parallel Graphs* (Centrum voor Wiskunde en Informatica, Amsterdam).

Sheffi Y (1985) *Urban Transportation Networks: Equilibrium Analysis With Mathematical Programming Methods* (Prentice-Hall, Englewood Cliffs, NJ).

Smith MJ (1979) The existence, uniqueness and stability of traffic equilibria. *Transportation Res. Part B: Methodological* 13(4): 295–304.

Steinberg R, Zangwill WI (1983) The prevalence of Braess' paradox. *Transportation Sci.* 17(3):301–318.

Valdes J, Tarjan RE, Lawler EL (1979) The recognition of series parallel digraphs. *Proc. 11th Annual ACM Sympos. Theory Comput.* (ACM, New York), 1–12.

Vetta A (2002) Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. *Proc. 43rd Annual IEEE Sympos. Foundations Comput. Sci.* (IEEE Computer Society, Los Alamitos, CA), 416–425.

Voorneveld M, Borm P, Van Megen F, Tijs S, Facchini G (1999) Congestion games and potentials reconsidered. *Internat. Game Theory Rev.* 1(3-4):283–299.

Wardrop JG (1952) Some theoretical aspects of road traffic research. *Proc. Instit. of Civil Engineers* 1(3):325–378.

Yang S, Hajek B (2005) Revenue and stability of a mechanism for efficient allocation of a divisible good. Working paper, University of Illinois at Urbana–Champaign, Champaign.

Daron Acemoglu is the Elizabeth and James Killian Professor of Economics in the Department of Economics at MIT, a member of the Institutions, Organizations and Growth program of the Canadian Institute for Advanced Research, and a research associate of the National Bureau Economic Research. He is the recipient of the John Bates Clark Medal in 2005, the 2012 Erwin Plein Nemmers Prize in Economics, the 2016 BBVA Frontiers of Knowledge Award in Economics, Finance and Management, and the Carnegie Fellowship in 2017.

Ali Makhdoumi is a Ph.D. candidate in the Electrical Engineering and Computer Science Department of MIT and a member of the Laboratory for Information and Decision Systems at MIT. His research interests include optimization, learning theory, game theory, and network science with applications to operations research.

Azarakhsh Malekian is an assistant professor of operations management and statistics at Rotman School of Management, University of Toronto. She is also a research affiliate at the Laboratory of Information and Decision Systems. Her research interests include game theory, network economics, and market design with applications to social and economic networks, online markets, and the sharing economy.

Asuman Ozdaglar is the Joseph F. and Nancy P. Keithley Professor and head of the Electrical Engineering and Computer Science Department at MIT. She is recipient of a Microsoft fellowship, the MIT Graduate Student Council Teaching Award, the NSF CAREER award, the 2008 Donald P. Eckman Award of the American Automatic Control Council, the Class of 1943 Career Development Chair, the inaugural Steven and Renee Innovation Fellowship, and the 2014 Spira teaching award.