MATHEMATICS OF OPERATIONS RESEARCH
Vol. 30, No. 1, February 2005, pp. 225–244
ISSN 0364-765X | EISSN 1526-5471 | 05 | 3001 | 0225

DOI 10.1287/moor.1040.0122

Topological Conditions for Uniqueness of Equilibrium in Networks

Igal Milchtaich
Department of Economics, Bar-Ilan University, 52900 Ramat-Gan, Israel,
milchti@mail.biu.ac.il, http://faculty.biu.ac.il/~milchti

Equilibrium flow in a physical network with a large number of users (e.g., transportation, communication, and computer networks) need not be unique if the costs of the network elements are not the same for all users. Such differences among users may arise if they are not equally affected by congestion or have different intrinsic preferences. Whether or not, for *all* assignments of strictly increasing cost functions, each user's equilibrium cost is the same in all Nash equilibria can be determined from the network topology. Specifically, this paper shows that in a two-terminal network, the equilibrium costs are always unique if and only if the network is one of several simple networks or consists of several such networks connected in series. The complementary class of all two-terminal networks with multiple equilibrium costs for *some* assignment of (user-specific) strictly increasing cost functions is similarly characterized by an embedded network of a particular simple type.

Key words: congestion externalities; nonatomic games; transportation networks; network topology; uniqueness of equilibrium
MSC2000 subject classification: Primary: 90B10; secondary: 91A13
OR/MS subject classification: Primary: Networks/graphs; secondary: games/group decisions, noncooperative
History: Received August 28, 2003; revised March 31, 2004.

1. Introduction. Different kinds of networks, such as transportation, communication, and computer networks exhibit congestion effects, whereby increased demand for certain network elements (e.g., roads, telecommunication lines, and servers) tends to downgrade their performance or increase the cost of using them. In such networks, the users' decisions (e.g., choice of routes) are interdependent in that their optimal choices (e.g., the fastest routes) depend on what the others do. If they all choose optimally, given the others' choices, then the users' choices constitute a Nash equilibrium. Even if the users are identical in all respects, due to the congestion externalities, their choices at equilibrium may differ. However, if the number of users is very large and each of them has a negligibly small effect on the others, then they have equal equilibrium payoffs or costs. Moreover, the payoffs or costs in *different* equilibria are the same (Aashtiani and Magnanti [1]). With a *heterogeneous* population of users (i.e., a multiclass network; Dafermos [5]), this need not be so. As the following example shows, if the users are not identical, and are differently affected by congestion, equilibrium costs may vary not only across users, but also from one Nash equilibrium to another.

EXAMPLE 1.1. A continuum of three classes of users travels on the two-terminal network shown in Figure 2(a) below. Each user has to choose one of the four routes connecting the users' common point of origin o and the common destination d. The cost of each route is the sum of the costs of its edges. For each user class, the cost of each edge e is given by an increasing affine function of the fraction x of the total population with a route that includes e. The fraction of the population in each class and the corresponding cost functions are given in Figure 1, where blank cells indicate prohibitively high costs. Clearly, users in each class effectively have a choice of only two routes from o to d: e_1 and $e_2 e_5$ for Class I users, $e_2 e_5$ and e_3 for Class II, and e_3 and $e_4 e_5$ for Class III. (The costs of the other two routes are prohibitively high.) If all the users choose the *first-mentioned* route for their class, their choices constitute a *strict* Nash equilibrium in that each user's cost is strictly less than it would be on the alternative route. The same is true if everyone chooses the *second-mentioned* route. However, the costs in the first equilibrium (≅3.41, 0.77, and 2.46 for Class I, II, and III users, respectively) are different from those in the second equilibrium

User class	Fraction of population	Cost functions				
		e_1	e_2	e_3	e_4	e_5
I	4/13	$3.1+x$	$8x$			x
II	5/13		x	$0.5+x$		x
III	4/13			$8x$	$2.1+x$	x

FIGURE 1. Table for Example 1.1.

(≅3.08, 0.88, and 3.02, respectively), and similarly for the mean cost (≅2.10 in the first equilibrium and 2.22 in the second).

In Example 1.1, neither of the two Nash equilibria Pareto dominates the other: For Class I users, the first equilibrium cost is higher, and for Class II and III, the second. In fact, for the network in Figure 2(a), this would be so for *any* assignment of cost functions. This is because the routes in this network, as well as in the essentially identical one 2(b), are linearly independent in the sense that each of them has an edge that is not in any other route. As shown by the author in Milchtaich [13, Theorem 3], this topological property implies that, for any assignment of cost functions, all the Nash equilibria are Pareto efficient. In the other two networks in Figure 2, the routes are *not* linearly independent. In these networks, some Nash equilibria may be strictly Pareto dominated by others.

EXAMPLE 1.2. A continuum of three classes of users travels from o to d on the network in Figure 2(c). The fraction of the population in each user class and the corresponding cost functions are given in Figure 3, where blank cells indicate prohibitively high costs. Each

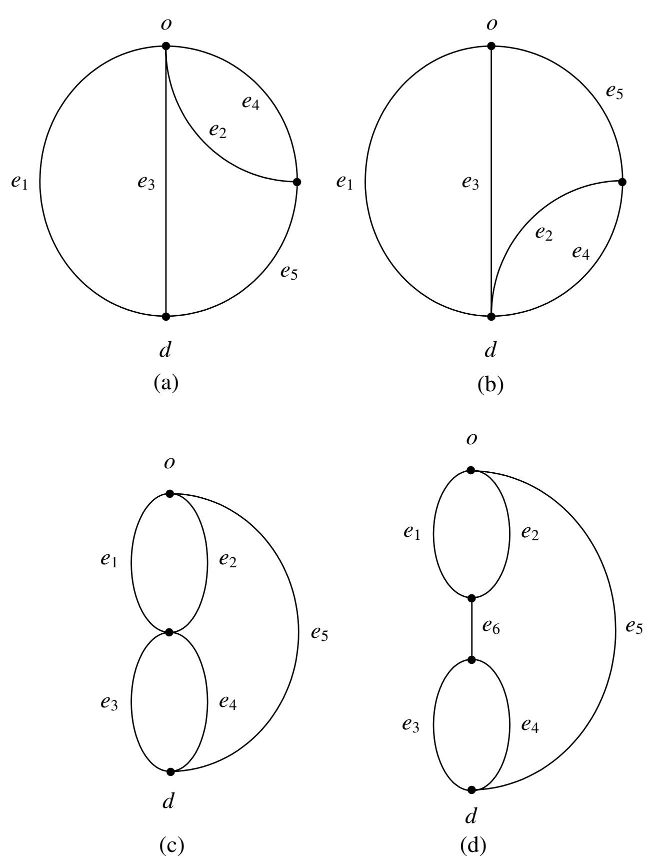

FIGURE 2. The *forbidden networks*: Two-terminal networks allowing for multiple equilibrium costs.

User class	Fraction of population	Cost functions				
		e_1	e_2	e_3	e_4	e_5
I	$3/7$		$6/7+3x$	$3x$		$6x$
II	$2/7$		$8x$		x	$6x$
III	$2/7$	x		$8x$		$6x$

FIGURE 3. Table for Example 1.2.

user can effectively choose only between e_5 and a single alternative route, which is $e_2\,e_3$ for Class I users, $e_2\,e_4$ for Class II, and $e_1\,e_3$ for Class III. There is one Nash equilibrium in which Class I users take e_5, and Class II and III take the respective alternative routes. In this equilibrium (which can be shown to be Pareto efficient), each user's cost is the same, $18/7$. There is another Nash equilibrium in which Class I users take their alternative route $e_2\,e_3$, and Class II and III take e_5. Again, the equilibrium cost is the same for all users, but this time it is higher, and equals $24/7$. Both equilibria are *not* strict: Any user taking e_5 would incur the same cost on his alternative route, and vice versa. Because the cost functions are affine, this implies that any convex combination of the two equilibria (in terms of the percentage of users in each class taking each route) is also a Nash equilibrium, with costs given by the corresponding convex combination of the above costs. Thus, there is a continuum of Nash equilibria, which can be Pareto ranked because, in each equilibrium, the costs for all users are the same.

The main result of this paper is that, if the costs to users are allowed to differ, whether or not there exist *some* cost functions with multiple equilibrium costs for some users depends on the network topology. The author (Milchtaich [12]) showed that in a network with parallel routes (like the one in Figure 4(a)), the equilibrium costs are unique for *any* assignment of cost functions. In this paper, it is shown that, in fact, the same is true for all five networks

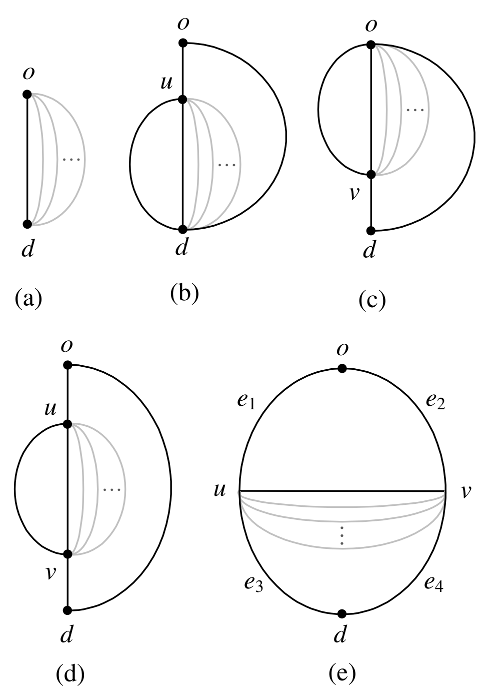

FIGURE 4. Two-terminal networks in which the equilibrium costs are always unique. The network in (a), in which one or more edges are connected in parallel (gray curves indicate optional edges), is embedded in each of the other four. Any network homeomorphic to one of these five networks is said to be *nearly parallel*.

in Figure 4, as well as for all the networks obtained by connecting two or more of them in series. Moreover, these are essentially the *only* two-terminal networks in which uniqueness of each user's equilibrium cost is guaranteed. For any other two-terminal network, it is possible to find an example with multiple equilibrium costs, very similar to Examples 1.1 or 1.2 above. Indeed, any such network has one of those in Figure 2 embedded in it, in a sense made precise below. These four networks, which will be referred to as the *forbidden networks*, are the *minimal* networks for which multiple equilibrium costs are possible. Thus, this paper gives two equivalent topological characterizations of two-terminal networks that may or may not have a multiplicity of equilibrium costs. The first directly identifies all networks with unique equilibrium costs for *any* assignment of strictly increasing cost functions, and the second all the networks in which, for *some* assignment of such cost functions, the equilibrium costs are not unique. Moreover, the results below show that, in the first kind of networks, not only are the equilibrium costs unique, but also the percentage of each class of users traversing each edge at equilibrium is *generically* unique. This entails that, unless certain special relations exist among the cost functions, this percentage is the same in all Nash equilibria.

In this paper, networks are always assumed to be undirected, in contrast to the more common practice in the literature of assuming that edges are *directed*, and can be traversed in one direction only. Here, such traveling restrictions, if they exist, are considered part of the *cost functions*, which may (but do not have to) assign a very high cost to one of the two directions. The merits of this approach are demonstrated by the *results* in this paper (and in Milchtaich [13]). These results show that the uniqueness of the equilibria is, indeed, linked to the topology of the underlying *undirected* network (and the same is true for their Pareto efficiency).

2. Graph-theoretic preliminaries. An *undirected multigraph* consists of a finite set of *vertices* together with a finite set of *edges*. Each edge e joins two distinct vertices, u and v, which are referred to as the *end vertices* of e. Thus, loops are not allowed, but more than one edge can join two vertices. An edge e and a vertex v are said to be *incident* with each other if v is an end vertex of e. The *degree* of a vertex is the number of edges incident with it. A *path* of length n is an alternating sequence p of vertices and edges $v_0e_1v_1 \cdots v_{n-1}e_nv_n$, beginning and ending with vertices, in which each edge is incident with the two vertices immediately preceding and following it and all the vertices (and necessarily all the edges) are distinct. Because of the latter assumption, each vertex and each edge in p either *precedes* or *follows* each of the other vertices and edges. The first and last vertices, v_0 and v_n, are called the *initial* and *terminal* vertices in p, respectively. If they are clear from the context, the path may be written more simply as $e_1e_2 \cdots e_n$. An *arc* is a path of length one, consisting of a single edge and its two end vertices. It may be viewed as an assignment of a particular *direction* to the edge. Obviously, each edge can be directed in two ways that differ from each other in the identity of the end vertex chosen as the initial vertex and that chosen as the terminal vertex. One, and only one, of these two arcs is contained in (i.e., is a section of) any path that includes the edge. In this sense, any such path specifies a particular direction for the edge. The set of all arcs in a network is denoted by $\mathcal{A}$.

A *two-terminal network* (*network*, for short) is an undirected multigraph together with a distinguished ordered pair of distinct vertices, o and d (for "origin" and "destination"), such that each vertex and each edge belong to at least one path with the initial vertex o and terminal vertex d. Any path r with these initial and terminal vertices will be called a *route*. The set of all routes in a network is denoted by $\mathcal{R}$.

Two networks G' and G'' will be identified if they are *isomorphic* in the sense that there is a one-to-one correspondence between the vertices of G' and G'' and between their edges, such that (i) the incidence relation is preserved, and (ii) the origin and destination in G' are paired with the origin and destination in G'', respectively. A network G' is a *subnetwork* of a network G'' if the former can be obtained from the latter by deleting a subset of its edges and nonterminal vertices (i.e., vertices other than o and d). A network G' will be said to be *embedded in the wide sense* in a network G'' if the latter can be obtained from the former by applying the following operations any number of times in any order (see Figure 5):

1. The *subdivision of an edge:* its replacement by two edges with a single common end vertex.
2. The *addition of a new edge* joining two existing vertices.
3. The *subdivision of a terminal vertex*, o or d: addition of a new edge e joining the terminal vertex with a new vertex v, followed by replacement of the former by the latter as the end vertex in two or more edges originally incident with the terminal vertex.

Two networks that can be obtained from the same network by successive subdivision of edges are said to be *homeomorphic*. Such networks can be obtained from each other by the insertion and removal of nonterminal degree-two vertices. For present purposes, they are nearly identical. By combining addition and subdivision of edges, complete new paths, which only have their initial and terminal vertices in a given network G, can be added to it. This is done by first joining the two vertices by a new edge e, and then subdividing e one or more times. Because any vertex and any edge in a network are in some route, this shows that if G is a subnetwork of another network, it is also embedded in it in the wide sense. A special case of terminal subdivision is *terminal extension*. In this operation, the new vertex v replaces the terminal vertex (o or d) as the end vertex of *all* the edges originally incident with the latter. The qualifier "in the wide sense" used in this paper is meant to distinguish the present notion of embedding from the more restrictive one in Milchtaich [13], which does not allow general terminal subdivisions but only terminal extensions. Terminal subdivisions may be viewed as selective terminal extensions: the new edge is appended to two or more, but not necessarily all, of the routes.

Two networks G' and G'' with the same origin-destination pair, but no other common vertices or edges, may be connected *in parallel*. The set of vertices in the resulting network G is the union of the sets of vertices in G' and G'', and similarly for the set of edges. The origin and destination in G are the same as in G' and G''. Two networks G' and G'' with a single common vertex (and, hence, without common edges), which is the destination in G' and the origin in G'', may be connected *in series*. The set of vertices in the resulting network G is the union of the sets of vertices in G' and G'', and similarly for the set of

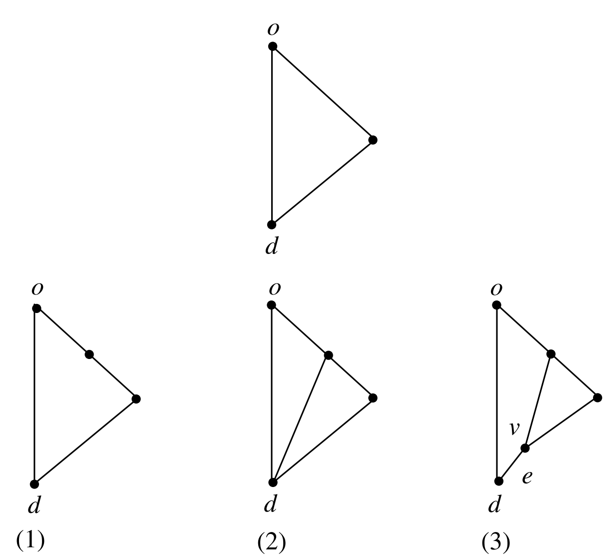

FIGURE 5. Embedding. The upper network is embedded in the wide sense in each of the lower three, which are obtained from it by (1) subdividing an existing edge, (2) adding a new edge, and, finally, (3) subdividing the destination.

edges. The origin in G coincides with the origin in G' and the destination is the destination in G''. The connection of an arbitrary number of networks in series or in parallel is defined recursively. Each of the connected networks is embedded in the wide sense in the network resulting from their connection.

A *parallel network* is one consisting of one or more edges connected in parallel, as in Figure 4(a). A network G will be said to be *nearly parallel* if it is homeomorphic to one of the five networks in Figure 4. It is not difficult to see that this is the case if and only if G has a single route, two parallel routes, or can be constructed from a network with two parallel routes by adding to it one or more parallel paths, with common initial and terminal vertices. The following graph-theoretic result plays an important role in this paper. The proof of this and the other results in the paper are given in §7.

Proposition 2.1. *For every two-terminal network G, one, and only one, of the following conditions holds*:

(i) *G is nearly parallel, or it consists of two or more nearly parallel networks connected in series.*

(ii) *One* (*or more*) *of the forbidden networks is embedded in the wide sense in G*.

3. The model. The population of *users* is modelled as an infinite set I (e.g., the unit interval $[0, 1]$), endowed with a nonatomic probability measure μ (e.g., Lebesgue measure), the *population measure*. This measure assigns values between zero and one to a σ-algebra of subsets of I, the *measurable sets*, which are interpreted as the set sizes relative to the total population. For a given network, a *strategy profile* is a mapping $\sigma\colon I \to \mathcal{R}$ (from users to routes) such that, for each route r in the network, the set of all users i with $\sigma(i) = r$ is measurable. For each arc a, the measure of the set of all users i such that a is contained in $\sigma(i)$ is called the *flow* on a in σ, and is denoted by f_a. Note that each edge e is associated with a *pair* of arc flows: one giving the flow on e in one direction and the other in the opposite direction. However, if all the routes in the network pass though e in the same direction, then one of these flows is always zero, and in this case, there is no ambiguity in associating e with a single arc flow, which may be denoted by f_e.

The cost of each arc a for each user i is given by a nonnegative and strictly increasing *cost function* $c_a^i\colon [0, 1] \to [0, \infty)$. When the flow on arc a is f_a, the cost for user i of traversing a is $c_a^i(f_a)$. Note that, for each user, each edge e is associated with a pair of cost functions, one for each direction. In each direction, the cost only depends on the flow on e in that direction, and not on the flow in the opposite direction. Although this assumption may involve some loss of generality, there are two important cases in which it does not. In the first case, all the routes in the network pass through e in the same direction, and, therefore, the flow in the opposite direction is always zero. Hence, only one cost function has to be associated with e for each user i, which, without ambiguity, may be denoted by c_e^i. In the other case, there is one direction of e which is prohibitively costly for all users. This would be true for *all* edges if passing through each edge were only allowed in one direction—an assumption made in much of the literature (e.g., Newell [15], Sheffi [18], Bell and Iida [4], and Nagurney [14]; but not Beckmann et al. [3]). For the sake of generality and simplicity of notation, in this paper the cost functions are *not* required to have the property that traversing an edge in a particular direction is very costly for all users. However, this requirement would not affect any of the results below, as long as there are no restrictions on *which* of the two directions has this property. The case of *predetermined* directionality is discussed in §6. Another assumption implicit in the definition of cost function is that the cost of each arc a for each user only depends on the total flow on a, and not on the *identities* of the other users. This does not imply that the congestion impacts of any two users are assumed to be equal, but rather that they are in fixed *proportion* to one another. Thus, for example, if one user (a bus, say) has twice the impact of another user (a sedan) in one arc, then the impact is also twofold in any other arc. The "size" of each group of users, given by the population measure, expresses its potential contribution to congestion, and, in heterogeneous populations, does not necessarily reflect the number of members.

The cost of each route r for each user i is defined as the sum of the costs for user i of the arcs contained in r. The cost thus depends on the flow on each of r's edges in the direction specified by the route. A strategy profile σ is a (pure-strategy) *Nash equilibrium* (in the nonatomic game defined by the network G and the cost functions) if each route is only used by those for whom it is a minimal-cost route. Formally, the equilibrium condition is: For each user i,

$$\sum_{\substack{a\in\mathcal{A}\\ \sigma(i)\text{ contains }a}} c_a^i(f_a) = \min_{r\in\mathcal{R}} \sum_{\substack{a\in\mathcal{A}\\ r\text{ contains }a}} c_a^i(f_a), \tag{1}$$

where, for each arc a, f_a is the flow on a in σ. For an equilibrium σ, the minimum in (1) is user i's *equilibrium cost*. In the special case of a transportation network with identical users, the above definition essentially reduces to the principle, formulated by Wardrop [19] and others (see Nagurney [14, p. 151]), that, at equilibrium, the travel time on all used routes is equal, and less than or equal to that of a single vehicle on any unused route.

4. Existence and uniqueness of equilibrium. Under weak assumptions on the cost functions, at least one Nash equilibrium exists. Specifically, a sufficient condition for the *existence* of equilibrium is that, for all arcs a, $c_a^i(x)$ is a continuous function of x for each user i and a measurable function of i for each $0 \le x \le 1$. The proof of this assertion, which is similar to that of Theorem 3.1 in Milchtaich [12], is omitted. The assertion can also be deduced from more general results, e.g., Schmeidler [17, Theorem 1] or Rath [16, Theorems 1 and 2]. The main concern of this paper is *uniqueness*. If all the users are identical, then the equilibrium itself is typically not unique. This is because, at equilibrium, any two groups of users of equal size taking different routes may interchange their choice of routes without affecting the equilibrium. However, the *equilibrium flow* on each arc in the network is the same in all Nash equilibria, which implies that the *equilibrium cost* is also unique (Aashtiani and Magnanti [1]). In fact, as the following proposition shows, this result extends to the case in which the users' cost functions are only *identical up to additive constants* in the sense that, for each pair of users i and i' and each arc a, the difference $c_a^i(x) - c_a^{i'}(x)$ is a constant that does not depend on x. For further extension of the uniqueness result, see Altman and Kameda [2].

Proposition 4.1. *If the users' cost functions are identical up to additive constants, then, for every two Nash equilibria, the flow on each arc in the network in the first equilibrium is equal to that in the second, and the same is also true for each user's equilibrium cost.*

For identical users (and, more generally, cost functions that are identical up to additive constants), uniqueness of the equilibrium costs would be guaranteed even if the cost functions were only assumed to be nondecreasing, rather than strictly increasing. However, for a heterogeneous population of users, the assumption of strict monotonicity (which in this paper is part of the definition of cost function) is critical. For example, suppose that one user class is totally unaffected by congestion, while another class is affected by it. If there are several minimal-cost routes for the first class of users, they may choose among them arbitrarily. Their choices do not affect their own equilibrium costs, but may affect those of the *second* user class. Therefore, regardless of the network topology, the equilibrium costs need not be unique.

The main question this paper addresses is whether, for a given network, the equilibrium arc flows and the equilibrium costs are unique for arbitrary (i.e., not necessarily identical up to additive constants) cost functions. A network will be said to have the (topological) *uniqueness property* if, for any assignment of (strictly increasing) cost functions, the flow on each arc is the same in all Nash equilibria. As the following proposition shows, this property can also be defined in terms of the equilibrium costs.

PROPOSITION 4.2. *For every two-terminal network G, the following three conditions are equivalent*:

(i) *G has the uniqueness property.*

(ii) *For any assignment of cost functions, each user's equilibrium cost is the same in all Nash equilibria.*

(iii) *For any assignment of cost functions, and for any pair of* strict *Nash equilibria,* some *user's equilibrium cost is the same in both equilibria.*

As mentioned in the introduction, whether the uniqueness property holds for a given network depends on its topology. Clearly, connecting two or more networks with the uniqueness property in series results in a network that also has this property, because the users' choice of routes in each constituent network does not restrict the choices or affect the costs in the other networks. In Milchtaich [12] and Konishi [8], the uniqueness property is shown to hold for all parallel networks. In fact, as the following theorem shows, this property holds for *all* five networks in Figure 4. Moreover, these networks and those constructed by connecting several of them in series are essentially the *only* two-terminal networks with the uniqueness property.

THEOREM 4.1. *A two-terminal network has the uniqueness property if and only if it is nearly parallel or it consists of two or more nearly parallel networks connected in series.*

An immediate corollary of Theorem 4.1 and Proposition 2.1 is that the networks *without* the uniqueness property are precisely those in which one of the forbidden networks is embedded in the wide sense. In this sense, these four networks are the *minimal* networks without this property. This result closely resembles Kuratowski's characterization of nonplanar graphs in terms of embedded forbidden graphs (Harary [7], Diestel [6]). In both cases, the set of minimal graphs or networks without a particular property (planarity, the uniqueness property) is finite (with two and four elements, respectively). In the case of planer graphs, this may be viewed as a corollary of Robertson and Seymour's minor graph theorem (see Diestel [6]). Whether a version of this theorem also holds for the present case of two-terminal networks and embedding in the wide sense is unknown.

COROLLARY 4.1. *For every two-terminal network G, there exists* some *assignment of cost functions with multiple equilibrium costs if and only if one of the forbidden networks is embedded in the wide sense in G.*

For the network in Figure 2(a), Example 1.1 specifies cost functions with two strict Nash equilibria such that each user's equilibrium cost is *not* the same in both equilibria. These cost functions give rise to the same equilibrium costs in the second network in Figure 2, which differs from the first only in that the origin and destination are interchanged. For the network in Figure 2(c), Example 1.2 specifies cost functions with two Nash equilibria, the first strictly Pareto dominating the second. In this example, both equilibria are not strict. However, it is easy to modify Example 1.2 such that the two equilibria become strict. For example, if the two cost functions of the form $8x$ are changed to $8.7x$, those of the form $6x$ to $6.5x$, and the constant $6/7$ to 1.1, each user's equilibrium cost is still lower in the first equilibrium than in the second, but in both equilibria, the cost of the user's equilibrium route is strictly less than that of the other routes. It follows that the modified cost functions can also be used for the network in Figure 2(d): Two strict Nash equilibria with arbitrarily close costs to those in the previous network can be obtained simply by assigning a sufficiently low cost (e.g., a cost function of $x/50$) to edge e_6.

User class	Fraction of population	Cost functions					
		e_1	e_2	e_3	e_4	e_5	e_6
I	9/20	x	$8x$	$8x$	$3x$	x	x
II	9/20	$8x$	$3x$	x	$8x$	x	x
III	1/10	$8x$	x	$8x$	x	x	$2x$

FIGURE 6. Table for Example 4.1.

Another assignment of cost functions with a multiplicity (indeed, continuum) of equilibrium costs for the network in Figure 2(d) is given in the next example. Note the very simple form of the cost functions in this example: linear, without constant terms. In Examples 1.1 and 1.2, differences among the users involve both the slopes of the cost curves, which reflect the degree to which different users are affected by congestion in each edge, and the intercepts, which represent their intrinsic preferences. (For example, some drivers may prefer shorter routes, even when the traffic flow on them is relatively slow, while for others, distance may be relatively less important than time.) Proposition 4.1 shows that differences of the latter kind are not sufficient for multiplicity of the equilibrium costs. As the following example shows, they are also not necessary.

EXAMPLE 4.1. A continuum of three classes of users travels from o to d on the network in Figure 2(d). The fraction of the population in each user class and the corresponding cost functions are given in Figure 6. There is one Nash equilibrium in which 5/18 of Class I users (1/8 of the total population) take the route $e_1 e_6 e_4$; 5/18 of Class II users take the route $e_2 e_6 e_3$; and all the other users take e_5. In this equilibrium, each user's equilibrium cost is the same, 0.75. There is another Nash equilibrium in which 5/36 of Class I users (1/16 of the total population) take the route $e_1 e_6 e_4$; 5/36 of Class II users take the route $e_2 e_6 e_3$; the rest of Class I and II users take e_5; and all Class III users take the route $e_2 e_6 e_4$. In this equilibrium, the cost for all users is 0.775.

By using these examples, it is possible to construct an assignment of cost functions with multiple equilibrium costs for *any* given network that does not have the uniqueness property. Of course, other such assignments also exist. Konishi [8] gives an example of such an assignment for a network in which the one in Figure 2(b) is embedded in the wide sense. Incidentally, with a *finite* number of nonidentical users (of equal mass, or congestion impact), the equilibrium costs need not be unique even in a parallel network (cf. Milchtaich [11]). The finite case also differs from that of a continuum of users in that an equilibrium need not *exist*. Although a Nash equilibrium in pure strategies always exists in a parallel network (Milchtaich [10, 11]), this is not so for the network obtained by adding a single edge joining o and d to the forbidden network in Figure 2(c) (with three players), or to the nearly parallel one in Figure 4(e) (with two players). Konishi [8] gives another example. The question of what topological conditions are both necessary and sufficient for the existence of equilibrium in the finite case is open.

5. Equivalence of equilibria. The uniqueness result for a heterogeneous population of users (Theorem 4.1) can be taken one step further. In a network with the uniqueness property, not only is the flow on each arc the same in all Nash equilibria, but *generically*, it is also made up of the same mixture of *user types*. While it is very easy, even for networks with the uniqueness property, to construct examples in which different types of users traverse a given arc in different equilibria, the generic uniqueness result entails that such examples depend on the existence of certain special relations among cost functions. If there is no a priori reason to assume that such relations exist, a unique composition of user types would be expected for each arc.

The theorem below is an extension of a similar result for parallel networks (Milchtaich [12, Theorem 4.3]). It is based on a very similar model to that used in that special case.

A *partition* of the population is a finite disjoint family of measurable sets $I_1, I_2, \ldots, I_n$, the *user classes*, such that $\bigcup_m I_m = I$ and $\mu(I_m) > 0$ for all m. (Extension to the case of an infinite—even uncountably infinite—family of user classes is possible. The formulation would closely follow that in Milchtaich [12].) A user class is interpreted as a group of users who are known to be of the same type. For given partition of the population and network G, let $\mathscr{G}$ denote the set of all assignments of continuous and strictly increasing cost functions c_a^i: $[0, 1] \to [0, \infty)$ with the property that, for every pair of users i and i' in the same class I_m, $c_a^i = c_a^{i'}$ for all arcs a. Because this property clearly implies that, for any fixed $0 \le x \le 1$, the mapping $i \mapsto c_a^i(x)$ is measurable, it follows from the remarks at the beginning of §4 that every element of $\mathscr{G}$ has a nonempty set of Nash equilibria. Two Nash equilibria σ and τ will be said to be *equivalent* if the contribution of each user class I_m to the flow on each arc a is the same in σ and τ; i.e., the measure of the set of all users $i \in I_m$ such that $\sigma(i)$ contains a is equal to the measure of the set of all users $i \in I_m$ such that $\tau(i)$ contains a. This condition clearly implies that the flow on each arc is the same in both equilibria. The *distance* between two elements of $\mathscr{G}$, one with cost functions c_a^i and the other with $\hat{c}_a^i$, is defined as $\max |c_a^i(x) - \hat{c}_a^i(x)|$, where the maximum is taken over all users i, arcs a, and $0 \le x \le 1$. This defines a metric for $\mathscr{G}$ (which, as shown in Milchtaich [12], is equivalent to some *complete* metric for $\mathscr{G}$, in the sense that the two metric topologies are the same. In other words, the metric space $\mathscr{G}$ is *topologically complete*). In a metric space, a property is considered to be *generic* if it holds in an open dense set (Mas-Colell [9, §8.2]). The following theorem asserts that the property that all Nash equilibria are equivalent is generic if and only if the network satisfies condition (i) in Proposition 2.1.

Theorem 5.1. *For every two-terminal network G, the following two conditions are equivalent*:

(i) *G is nearly parallel, or it consists of two or more nearly parallel networks connected in series.*

(ii) *For every partition of the population, there is an open dense set in the space $\mathscr{G}$ such that, for any assignment of cost functions that belongs to this set, every two Nash equilibria are equivalent.*

Put in another way, condition (ii) states that, for every partition of the population, the set of all assignments of cost functions in $\mathscr{G}$ with two (or more) nonequivalent Nash equilibria is nowhere dense in $\mathscr{G}$. Together, Theorems 4.1 and 5.1 imply the following addition to Proposition 4.2.

Corollary 5.1. *Condition* (*ii*) *in Theorem* 5.1 *is equivalent to the uniqueness property.*

6. Remarks. The results in this paper, which link network topology with uniqueness of arc flows and equilibrium costs, are similar in spirit to those of Milchtaich [13], which link network topology with Pareto efficiency of equilibria. However, uniqueness and Pareto efficiency are each equivalent to a *different* topological property. Specifically, a two-terminal network has the property that, for any assignment of cost functions of the form considered here, all the equilibria are Pareto efficient if and only if it is a network with *linearly independent routes* in the sense that every route has at least one edge that is not in any other route (Milchtaich [13, Theorem 3]). On the other hand, linear independence of the routes is not a necessary or sufficient condition for the network to have the uniqueness property, because it holds for the first two networks in Figure 2, but not for the last two, and for the first four networks in Figure 4, but not for the last one.

A weaker topological property than linearly independent routes, which holds for all the networks in Figure 2, is a *series-parallel network*, i.e., one which can be constructed from single edges by sequentially connecting networks in series or in parallel. However, because the network in Figure 4(e) is not series-parallel, even this is not a necessary (or sufficient) condition for the network to have the uniqueness property. Theorem 1 in Milchtaich [13] shows that for a population of *identical* users, a series-parallel network is a necessary and sufficient condition for Braess's paradox never to occur. *Braess's paradox* is said to occur when lowering the cost of one or more arcs *increases* the users' equilibrium cost. With nonidentical users, a natural generalization of Braess's paradox can occur even in a series-parallel network (but not in one with linearly independent routes; see Milchtaich [13]). For example, replacing the constant $6/7$ in Example 1.2 with a higher value would only leave the first equilibrium mentioned, in which each user's equilibrium cost is $18/7$. Replacing it with any positive constant *less* than $6/7$ would leave only the second equilibrium, in which all the users have equilibrium costs *higher* than $18/7$.

This paper and Milchtaich [13] both consider undirected networks, and view *directionality*, if it exists, as part of the cost functions. In a series-parallel network, edges have intrinsic directions, because all routes pass through each edge in the same direction (see Milchtaich [13]). Of the networks in Figures 2 and 4, only that in Figure 4(e) is not series-parallel, and in this network, uniqueness of the equilibrium cost does not depend on how the edges joining u and v are directed. Nevertheless, the results in this paper would not hold if edges were viewed as having predetermined directions. This is demonstrated by the two directed networks in Figure 7. The undirected version of the network in Figure 7(a) is homeomorphic to that in Figure 4(e). Therefore, for any assignment of cost functions, the equilibrium flow on each arc is unique and, in addition, for any partition of the population, there is an open dense set in the space $\mathcal{G}$ such that, for every assignment of cost functions in this set, all Nash equilibria are equivalent. The same clearly holds for the second directed network in Figure 7, in which the directed routes are essentially the same as in the first. However, this directed network cannot be constructed by connecting in series directed versions of nearly parallel networks. This shows that Theorems 4.1 and 5.1 do not hold for directed networks. The directed network in Figure 7(b) also cannot be obtained by subdivision of edges, addition of edges, or subdivision of terminal vertices from any of the forbidden networks if these are directed in the natural manner as series-parallel networks. This shows that Proposition 2.1 also does not hold for directed networks, which demonstrates the usefulness of the present approach of linking uniqueness of equilibrium costs with the topology of the *undirected* network.

This paper only considers networks with two terminal vertices. In one sense, this limitation is not too severe. A network with multiple origin-destination pairs may be modelled by connecting all origins to a single, fictitious vertex o, and similarly for the destinations. The restriction of each user i to a specific origin can be implemented by assigning a very high cost to the edges joining o with the other origins, and similarly for i's destination. Thus, user-specific cost functions partially compensate for the topological limitation of a single origin-destination pair. If the augmented network, constructed as above, has the uniqueness property, then the equilibrium costs are unique also for any assignment of cost functions and origin-destination pairs in the original network. However, it is easy to see that the uniqueness property is not a *necessary* condition for this. Thus, user-specific cost functions cannot completely circumvent the topological limitation.

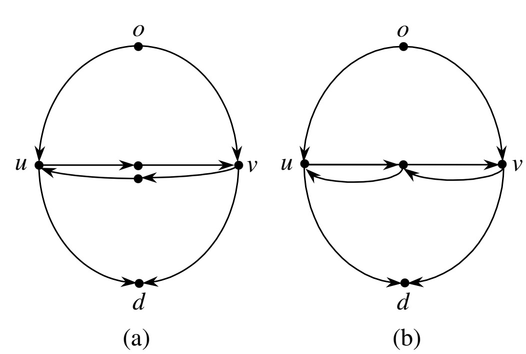

FIGURE 7. Two directed two-terminal networks in which the equilibrium costs are always unique.

7. Proofs. This section presents the proofs of the various results in the paper. The proof of Proposition 2.1 requires the following lemma.

LEMMA 7.1. *For every two-terminal network G, at least one of the two conditions in Proposition 2.1 holds.*

PROOF. The proof proceeds by induction on the number of edges in G. If there is only one edge, condition (i) in Proposition 2.1 clearly holds. Suppose that G has more than one edge. The induction hypothesis will be that the assertion of the lemma holds for every network with a smaller number of edges than G. If G can be constructed by connecting two other networks in series, then, by the induction hypothesis, (i) each of these networks is nearly parallel or consists of several nearly parallel networks connected in series, or (ii) a forbidden network is embedded in the wide sense in at least one of them. In the first case, condition (i) in Proposition 2.1 holds for G, and in the second, condition (ii) holds. Suppose next that G has more than one edge, but it cannot be constructed by connecting two other networks in series. By Milchtaich [13, Lemma 1] (which is essentially part of Harary [7, Theorem 3.3]), this implies that there are two routes in G that do not have any common edges or vertices other than o and d (i.e., parallel routes). The edges and vertices in these two routes constitute a subnetwork of G, which is homeomorphic to a two-edge parallel network. Consider the collection of *all* nearly parallel subnetworks of G with two or more routes. In this collection, choose a *maximal* subnetwork G', i.e., one that is not itself a subnetwork of any other member of the collection. If $G' = G$, then G is nearly parallel, and the proof is complete. Suppose, then, that some route r in G includes an edge e that is not in G'. Let u be the last vertex before e in r that *is* in G', and v the first vertex after e in r that is in G'. None of the edges and vertices in r that follow u and precede v are in G'. Adding these edges and vertices to G' results in a subnetwork of G that is homeomorphic to a network G'' obtained by adding a *single* edge to G'. The assumed maximality of G' implies that G'' is not nearly parallel. Therefore, to complete the proof of the lemma, it suffices to establish the following.

CLAIM. *Let G'' be a network that results from adding a single edge e to a nearly parallel network G' with two or more routes. Then, one of the forbidden networks is embedded in the wide sense in G'' or G'' is nearly parallel.*

The proof of the claim involves checking five cases, (a) through (e). In each case, G' is assumed to be homeomorphic to the corresponding network in Figure 4.

Case (a). If the end vertices of e (the edge added to G') are o and d, then G'' is nearly parallel (specifically, homeomorphic to a parallel network). If at least one end vertex is not o or d, then depending on whether the network G' has (i) only two or (ii) three or more routes, the network G'' is (i) homeomorphic to one of the networks in Figures 4(b)–4(e), or (ii) homeomorphic to a network obtained by adding one or more edges joining o and d to one of these four networks. In the first case, G'' is nearly parallel, and in the second, one of the networks in Figures 2(a) and 2(b) is embedded in it in the wide sense.

Case (b). In this case, there is a unique nonterminal vertex u in G' of degree three or higher, and a unique route r that does *not* include u. If both end vertices of e are in r (possibly coinciding with its initial or terminal vertices o or d), then the network in Figure 2(b) is embedded in the wide sense in G''. (For example, if one end vertex is o and the other, v, is *not* d, then obtaining G'' from the network in Figure 2(b) involves subdivision of d, possibly followed by a number of edge additions and subdivisions. The subdivision of d adds a new vertex v and an edge joining it with d, and replaces d by v as the end vertex of e_1 and e_3, but not of e_2 and e_4, which remain incident with d.) If only one of the two end vertices of e is in r, and this vertex is not d, then the other end vertex is in some route which also includes u. Depending on whether that vertex follows, coincides with, or precedes u, the network in Figure 2(b), 2(c), or 2(d), respectively, is embedded in the wide sense in G''. The same three possibilities exist if both end vertices of e are in routes that also include u, and at least one of them precedes it. If one end vertex coincides with u and the other one follows it, then depending on where the latter vertex lies, G'' is nearly parallel or has the network in Figure 2(a) embedded in it in the wide sense. Finally, if both end vertices of e follow u, then the network in Figure 2(b) is embedded in the wide sense in G''.

Case (c). This case is very similar to the previous one because the network in Figure 4(c) is obtained from that in 4(b) by interchanging o and d.

Case (d). In this case, there are two nonterminal vertices in G' of degree three or higher, u and v, and a unique route r that does not include either of them. The analysis of the present case is identical, verbatim, to that of Case (b) except for the final sentence, which has to be modified as follows. If both end vertices of e follow u, then depending on whether at least one of them precedes v, one of them coincides with v and the other one follows it, or both of them follow v, the network in Figure 2(b), 2(c), or 2(d), respectively, is embedded in the wide sense in G''. (For example, if one end vertex of e is d and the other one follows u but precedes v, then obtaining G'' from the network in Figure 2(b) involves: (i) subdivision of d, which adds the vertex v and an edge joining it with d, and replaces d by v as the end vertex of e_2 and e_3, but not of e_1 and e_4; followed by (ii) subdivision of o, which adds the vertex u and an edge joining it with o, and replaces o by u as the end vertex of e_3 and e_5, but not of e_1; possibly followed by (iii) a number of edge additions and subdivisions.)

Case (e). The network in Figure 4(e) can be obtained from that in 4(b) by subdivision of d and from that in 4(c) by subdivision of o. Therefore, it is not difficult to see that a network homeomorphic to G'' can be obtained from a network homeomorphic to either 4(b) or 4(c) by (i) addition of an edge, followed by (ii) subdivision of a terminal vertex. Clearly, if the network obtained in the interim stage (after the first operation is carried out) is also homeomorphic to one of the networks in Figures 4(b) and 4(c), then G'' is homeomorphic to that in Figure 4(e). If the network obtained in the interim stage is not homeomorphic to one of these two networks, then it follows from the analysis of Cases (b) and (c) that one of the forbidden networks is embedded in it in the wide sense, and in this case, the same is true for G''. □

Proof of Proposition 2.1. In view of Lemma 7.1, it suffices to show that if a network satisfies condition (i), then no forbidden network is embedded in it in the wide sense. Clearly, a forbidden network does not satisfy condition (i), because none of the networks in Figure 2 is homeomorphic to any of those in Figure 4, or can be obtained by connecting any two other networks in series. Therefore, it is sufficient to prove the following.

Claim. *Suppose that a network G' is embedded in the wide sense in a network G'' that satisfies condition* (i). *Then, G' also satisfies* (i).

To prove this claim, it clearly suffices to consider the case in which G'' is obtained from G' by one of the three operations defining embedding in the wide sense. If the operation is the subdivision of an edge, then the two networks are homeomorphic, and so one of them satisfies condition (i) if and only if the other network does so. Next, suppose that G'' is obtained from G' by the addition of a single edge e joining two existing vertices. Assume, without loss of generality, that G'' is nearly parallel. Both end vertices of e are terminal vertices or have degree three or higher. Hence, each of them can only be one of those marked o, d, u, or v in the various networks in Figure 4. It follows that G', which can be recovered by removing e from G'', is necessarily a nearly parallel network or can be obtained from one by connecting it in series with one or two networks with single edges. For example, if G'' is homeomorphic to the network in Figure 4(e), then G' must be homeomorphic to one of those in Figures 4(a) and 4(e) (this is the case if e joins u and v), or to a network obtained by connecting in series a network with a single edge and one of those in Figures 4(a)–4(c). This proves that, in the case of the addition of an edge, G' satisfies condition (i).

It remains to consider the case in which G'' is obtained from G' by the subdivision of a terminal vertex, say d. The vertex v and edge e thereby created belong to a nearly parallel network G, which (by condition (i)) either coincides with G'' or gives it after connection in series with another network satisfying (i). The latter possibility necessarily holds if e is the *only* edge incident with d, in which case the network connected in series with G is the original network G', which hence satisfies condition (i). Suppose, then, that e is not the only edge incident with d. Every other edge incident with d is included in some route that does *not* include v. This implies that v is a nonterminal vertex in G, and by definition of terminal subdivision, the degree of v is three or higher. These properties of e and v essentially identify them uniquely. In particular, they imply that G cannot be homeomorphic to one of the networks in Figures 4(a) and 4(b), in which a vertex and an edge with these properties do not exist. If G is homeomorphic to one of the networks in Figures 4(c) and 4(d), then v must be the vertex marked as such, and e the edge joining it with d. This shows that, before the subdivision of d, the network that is now G was homeomorphic to one of the networks in Figures 4(a) and 4(b), which proves that condition (i) holds for the original network G'. The same conclusion holds if G is homeomorphic to the network in Figure 4(e), in which case similar reasoning shows that the network that is now G was formerly homeomorphic to that in Figure 4(b).

An independent argument showing that a network cannot satisfy *both* conditions in Proposition 2.1 is given in the proof of Proposition 4.2 below. Together with Lemma 7.1, it constitutes an alternative, partially "game-theoretic," proof for Proposition 2.1.

Proof of Proposition 4.1. Suppose that the users' cost functions are identical up to additive constants, and let i_0 be one of the users. Then,

$$c_a^i(x) - c_a^i(y) = c_a^{i_0}(x) - c_a^{i_0}(y) \tag{2}$$

for all users i, arcs a, $0 \le x \le 1$ and $0 \le y \le 1$. Let σ and $\hat{\sigma}$ be two Nash equilibria, and, for each arc a, let f_a and $\hat{f}_a$ be the flow on a in σ and $\hat{\sigma}$, respectively. For every user i and route r, define $\sigma_r(i)$ as 1 or 0 according to whether $\sigma(i)$ is equal to or different from r, respectively. Define $\hat{\sigma}_r(i)$ in a similar manner. Because σ is a Nash equilibrium, it follows from (1) that, for all users i,

$$\sum_{r\in\mathcal{R}}\left[\sum_{a\in\mathcal{A},\ r \text{ contains } a} c_a^i(f_a)\right](\sigma_r(i) - \hat{\sigma}_r(i)) \le 0.$$

Because $\hat{\sigma}$ is a Nash equilibrium, similar inequalities hold with σ and $\hat{\sigma}$ interchanged and f_a replaced by $\hat{f}_a$. It follows that, for all users i,

$$\sum_{r\in\mathcal{R}}\left[\sum_{a\in\mathcal{A},\ r \text{ contains } a} (c_a^i(f_a) - c_a^i(\hat{f}_a))\right](\sigma_r(i) - \hat{\sigma}_r(i)) \le 0.$$

Changing the order of summation and using (2) gives that, for all users i,

$$\sum_{a\in\mathcal{A}}\left[\sum_{r\in\mathcal{R},\ r \text{ contains } a} (\sigma_r(i) - \hat{\sigma}_r(i))\right](c_a^{i_0}(f_a) - c_a^{i_0}(\hat{f}_a)) \le 0.$$

Integration over i now gives

$$\sum_{a\in\mathcal{A}}(f_a - \hat{f}_a)(c_a^{i_0}(f_a) - c_a^{i_0}(\hat{f}_a)) \le 0.$$

By strict monotonicity of the cost functions, each term in the last sum is nonnegative, and is moreover positive if $f_a \ne \hat{f}_a$. Therefore, all terms must be zero, and moreover, $f_a = \hat{f}_a$ must hold for all arcs a. This implies that the cost of each route for each user is the same in σ and $\hat{\sigma}$ and, therefore, the equilibrium costs are also equal. □

The following four lemmas are required for the proof of Proposition 4.2.

LEMMA 7.2. *The uniqueness property holds for a network G if and only if it holds for every network homeomorphic to it, and in this case, it also holds for every network obtained from G by removal of a single edge. If G can be constructed by connecting two other networks G' and G'' in series, then the uniqueness property holds for G if and only if it holds for both G' and G''.*

PROOF. To prove the first assertion, it clearly suffices to consider a network G' obtained from G by either the subdivision or removal of a single edge e. In the latter case, the uniqueness property clearly holds for G' if it holds for G, because removing an edge is equivalent to forcing its cost to be prohibitively high for all users. In the former case, only the *sum* of the costs of the two parts of e (in each of the two directions) matters, because any route in G' that includes one of them also includes the other. Therefore, for any given cost function for e in a particular direction, subdividing this edge and assigning half the original cost to each of its two parts has no effect on the costs of the routes. It is, therefore, clear that the uniqueness property holds for G' if and only if it holds for G.

If G results from connecting two networks G' and G'' in series, then there is a natural one-to-one correspondence between the set of all routes in G and the Cartesian product of the set of all routes in G' and the set of all routes in G''. This induces a one-to-one correspondence between the set of all strategy profiles in G and the Cartesian product of the set of all strategy profiles in G' and that of all strategy profiles in G''. There is also a one-to-one correspondence, defined by restrictions, between the set of all assignments of cost functions for G and the Cartesian product of the set of all assignments of cost functions for G' and those for G''. It is not difficult to see that a strategy profile σ in G is a Nash equilibrium with respect to a given assignment if and only if the corresponding strategy profiles σ' and σ'' in G' and G'' are Nash equilibria with respect to the respective restrictions. If G does not have the uniqueness property, then there is an assignment of cost functions for G with two Nash equilibria σ and τ such that the flow on some arc a in σ is different from that in τ. If a is in G', say, then these arc flows are the same as in σ' and τ', respectively, and therefore G' does not have the uniqueness property. Conversely, if G' does not have the uniqueness property, then any assignment of cost functions for G' with two Nash equilibria with different arc flows, and any assignment of cost functions for G'' with at least one Nash equilibrium, together define an assignment of cost functions for G for which the equilibrium arc flows are not unique. □

LEMMA 7.3. *The uniqueness property holds for all the networks in Figure 4.*

PROOF. To prove that the network in Figure 4(a) has the uniqueness property, it suffices to show that those in 4(b) and 4(c) have it. This is because, if the first network has more than one edge, connecting it in series with a network with a single edge gives a network that can also be obtained from that in Figure 4(b) or the one in 4(c) by removing the edge joining o and d. Therefore, it follows from Lemma 7.2 that if the networks in Figures 4(b) and 4(c) have the uniqueness property, so does that in 4(a). To prove that these two networks indeed have the uniqueness property, it suffices to show that the one in Figure 4(e) has it. This is because a network with a single edge connected in series with the network in Figure 4(b) or 4(c) can be obtained by the removal of one of the edges incident with the origin or the destination, respectively, in 4(e). Finally, the network in Figure 4(d) has the uniqueness property if and only if the one in 4(c) has it. This is because, in the former network, the edge preceding u and that following v together affect route costs in the same way as the single edge following v does in the latter network. In conclusion, it suffices to prove that the uniqueness property holds for the network in Figure 4(e).

Let σ and $\hat{\sigma}$ be two Nash equilibria with respect to the same assignment of cost functions for the network in Figure 4(e). For each arc a, let f_a be the flow on a in σ, $\hat{f}_a$ the flow in $\hat{\sigma}$, and $\Delta f_a = f_a - \hat{f}_a$. Clearly, $\Delta f_{e_1} + \Delta f_{e_2} = \Delta f_{e_3} + \Delta f_{e_4} = 0$. It has to be shown that $\Delta f_a = 0$ for all arcs a.

Claim 1. *If $\Delta f_{e_1} \geq 0 \geq \Delta f_{e_3}$, then $\Delta f_a = 0$ for all arcs a.*

Suppose that the assumption of the claim holds, or, equivalently, $\Delta f_{e_2} \leq 0 \leq \Delta f_{e_4}$. Let $\mathcal{A}_1$ be the set of all arcs with the initial vertex u and terminal vertex v, and $\mathcal{A}_2$ the set of all arcs with the initial vertex v and terminal vertex u. Each arc in $\mathcal{A}_1$ is contained in a unique route in G. Let $\mathcal{R}_1^+$ be the set of all routes that contain some arc $a \in \mathcal{A}_1$ with $\Delta f_a > 0$. Similarly, let $\mathcal{R}_2^-$ be the set of all routes that contain some arc $a \in \mathcal{A}_2$ with $\Delta f_a < 0$. For each user i, the cost of the equilibrium route $\sigma(i)$ is not greater than that of any alternative route. Therefore, if $\sigma(i)$ contains some arc $a \in \mathcal{A}_1$, necessarily $c^i_{e_1}(f_{e_1}) + c^i_a(f_a) \leq c^i_{e_2}(f_{e_2})$, $c^i_a(f_a) + c^i_{e_4}(f_{e_4}) \leq c^i_{e_3}(f_{e_3})$, and $c^i_a(f_a) \leq c^i_{a'}(f_{a'})$ for all $a' \in \mathcal{A}_1$. If, in addition, $\sigma(i) \in \mathcal{R}_1^+$ (i.e., $\Delta f_a > 0$), then it follows from the assumption $\Delta f_{e_1}, \Delta f_{e_4} \geq 0 \geq \Delta f_{e_2}, \Delta f_{e_3}$ that $c^i_{e_1}(\hat{f}_{e_1}) + c^i_a(\hat{f}_a) < c^i_{e_2}(\hat{f}_{e_2})$, $c^i_a(\hat{f}_a) + c^i_{e_4}(\hat{f}_{e_4}) < c^i_{e_3}(\hat{f}_{e_3})$, and $c^i_a(\hat{f}_a) < c^i_{a'}(\hat{f}_{a'})$ for all $a' \in \mathcal{A}_1$ with $\Delta f_{a'} \leq 0$. In this case, the route $\hat{\sigma}(i)$ *cannot* include the edges e_2 or e_3 or contain any arc $a' \in \mathcal{A}_1$ with $\Delta f_{a'} \leq 0$ (because less costly alternatives exist, and $\hat{\sigma}$ is an equilibrium). This proves that, for all users i with $\sigma(i) \in \mathcal{R}_1^+$, also $\hat{\sigma}(i) \in \mathcal{R}_1^+$. Therefore, if $\mathcal{R}_1^+$ is not empty, then there must be some $a \in \mathcal{A}_1$ with $\Delta f_a > 0$ such that the measure of the set of all users i for whom $\sigma(i)$ contains a is less than or equal to the measure of the set of users i for whom $\hat{\sigma}(i)$ contains a. However, this implies that $f_a \leq \hat{f}_a$, which contradicts the assumption $\Delta f_a > 0$. This contradiction proves that $\mathcal{R}_1^+$ is empty. A very similar argument shows that $\mathcal{R}_2^-$ is empty. It follows that the difference between the total flow on all the arcs belonging to $\mathcal{A}_1$ and that on all the arcs belonging to $\mathcal{A}_2$ (i.e., the net flow from u to v, which may be positive of negative) is either (i) the same in σ and $\hat{\sigma}$, or (ii) smaller in the former. In addition, (i) holds if and only if $\Delta f_a = 0$ for all arcs a belonging to either $\mathcal{A}_1$ or $\mathcal{A}_2$. It is not difficult to see that if (i) holds, then $\Delta f_{e_1} = \Delta f_{e_3}$, and if (ii) holds, then $\Delta f_{e_1} < \Delta f_{e_3}$. Because the last inequality contradicts the assumption $\Delta f_{e_1} \geq \Delta f_{e_3}$, this proves that $\Delta f_a = 0$ for *all* arcs a in the network.

Claim 2. *If $\Delta f_{e_1} > 0$, then $\Delta f_{e_3} \leq 0$.*

This will be proved by assuming that $\Delta f_{e_1}, \Delta f_{e_3} > 0$ (and, hence, $\Delta f_{e_2}, \Delta f_{e_4} < 0$), and showing that this assumption leads to a contradiction. Let I_σ be the set of all users i such that e_4 is in $\sigma(i)$ but not in $\hat{\sigma}(i)$, and $I_{\hat{\sigma}}$ the set of all users i such that e_4 is in $\hat{\sigma}(i)$, but not in $\sigma(i)$. The difference between the measures of these two sets, $\mu(I_\sigma) - \mu(I_{\hat{\sigma}})$, equals Δf_{e_4}. $\mathcal{A}_1$, $\mathcal{A}_2$, $\mathcal{R}_1^+$, and $\mathcal{R}_2^-$ are as in the proof of Claim 1.

Consider any user $i \in I_{\hat{\sigma}}$ such that $\hat{\sigma}(i)$ contains some arc $a \in \mathcal{A}_1$. Because $\hat{\sigma}$ is an equilibrium, $c^i_a(\hat{f}_a) + c^i_{e_4}(\hat{f}_{e_4}) \leq c^i_{e_3}(\hat{f}_{e_3})$. Because σ is an equilibrium and e_4 is not in $\sigma(i)$, $c^i_{e_3}(f_{e_3}) \leq c^i_a(f_a) + c^i_{e_4}(f_{e_4})$. It follows, by strict monotonicity and nonnegativity of the cost functions and the assumption $\Delta f_{e_3} > 0 > \Delta f_{e_4}$, that $c^i_a(\hat{f}_a) < c^i_a(f_a)$ and $c^i_{e_4}(f_{e_4}) < c^i_{e_3}(f_{e_3})$. Because e_4 is *not* in $\sigma(i)$, the latter inequality implies that $\sigma(i) = e_1 e_3$. The former inequality implies that $\Delta f_a > 0$, and, hence, $\hat{\sigma}(i) \in \mathcal{R}_1^+$. Next, consider any user $i \in I_{\hat{\sigma}}$ such that $\hat{\sigma}(i)$ does *not* contain any arc in $\mathcal{A}_1$, and, hence, $\hat{\sigma}(i) = e_2 e_4$. Because $\hat{\sigma}$ is an equilibrium, (i) $c^i_{e_2}(\hat{f}_{e_2}) + c^i_{e_4}(\hat{f}_{e_4}) \leq c^i_{e_1}(\hat{f}_{e_1}) + c^i_{e_3}(\hat{f}_{e_3})$, and (ii) $c^i_{e_4}(\hat{f}_{e_4}) \leq c^i_a(\hat{f}_a) + c^i_{e_3}(\hat{f}_{e_3})$ for all $a \in \mathcal{A}_2$. Because $\Delta f_{e_2}, \Delta f_{e_4} < 0 < \Delta f_{e_1}, \Delta f_{e_3}$, (i) implies that $c^i_{e_2}(f_{e_2}) + c^i_{e_4}(f_{e_4}) < c^i_{e_1}(f_{e_1}) + c^i_{e_3}(f_{e_3})$, from which it follows that $\sigma(i) \neq e_1 e_3$. Because by assumption $\sigma(i)$ does not include e_4, it must contain some $a \in \mathcal{A}_2$ for which $c^i_a(f_a) + c^i_{e_3}(f_{e_3}) \leq c^i_{e_4}(f_{e_4})$. Because $\Delta f_{e_3} > 0 > \Delta f_{e_4}$, $c^i_a(f_a) + c^i_{e_3}(\hat{f}_{e_3}) < c^i_{e_4}(\hat{f}_{e_4})$. It follows, by comparison with (ii) above, that $\Delta f_a < 0$, and, hence, $\sigma(i) \in \mathcal{R}_2^-$. Together, this and the previous conclusion prove that $I_{\hat{\sigma}}$ decomposes into two disjoint sets: the set $I_{\hat{\sigma}_1}$ of all users i with $\hat{\sigma}(i) \in \mathcal{R}_1^+$ and $\sigma(i) = e_1 e_3$, and the set $I_{\hat{\sigma}_2}$ of all users i with $\hat{\sigma}(i) = e_2 e_4$ and $\sigma(i) \in \mathcal{R}_2^-$. Hence, $\mu(I_{\hat{\sigma}}) = \mu(I_{\hat{\sigma}_1}) + \mu(I_{\hat{\sigma}_2})$. Let I_{σ_1} and I_{σ_2} be the subsets of I_σ defined in a similar manner to $I_{\hat{\sigma}_1}$ and $I_{\hat{\sigma}_2}$, but with σ and $\hat{\sigma}$ interchanged. Because these sets are clearly disjoint,

$\mu(I_{\sigma_1})+\mu(I_{\sigma_2}) \le \mu(I_\sigma)$. Therefore, $(\mu(I_{\sigma_1})-\mu(I_{\hat\sigma_1}))+(\mu(I_{\sigma_2})-\mu(I_{\hat\sigma_2})) \le \mu(I_\sigma)-\mu(I_{\hat\sigma}) = \Delta f_{e_4} < 0$, which implies that $\mu(I_{\sigma_1}) < \mu(I_{\hat\sigma_1})$ or $\mu(I_{\sigma_2}) < \mu(I_{\hat\sigma_2})$. However, as shown below, each of these two inequalities leads to a contradiction.

Suppose that $\mu(I_{\sigma_1}) < \mu(I_{\hat\sigma_1})$. This means that there are more users i with $\hat\sigma(i) \in \mathscr{R}_1^+$ and $\sigma(i) = e_1 e_3$ than users i with $\sigma(i) \in \mathscr{R}_1^+$ and $\hat\sigma(i) = e_1 e_3$. By definition of $\mathscr{R}_1^+$ (each element of which is the unique route containing a particular arc $a \in \mathscr{A}_1$ with $\Delta f_a > 0$), for every $r \in \mathscr{R}_1^+$, there are more users i with $\sigma(i) = r$ than users i with $\hat\sigma(i) = r$. Therefore, there are more users i with $\sigma(i) \in \mathscr{R}_1^+$ and $\hat\sigma(i) \notin \mathscr{R}_1^+$ than users i with $\hat\sigma(i) \in \mathscr{R}_1^+$ and $\sigma(i) \notin \mathscr{R}_1^+$. Because the latter kind of users includes all those for whom $\hat\sigma(i) \in \mathscr{R}_1^+$ and $\sigma(i) = e_1 e_3$, it follows from the assumption at the beginning of this paragraph that there are some users i with $\sigma(i) \in \mathscr{R}_1^+$ and $\hat\sigma(i) \notin \mathscr{R}_1^+$ for whom $\hat\sigma(i) \neq e_1 e_3$. For each of these users i, $\sigma(i)$ contains some arc $a \in \mathscr{A}_1$ with $\Delta f_a > 0$. Because σ is an equilibrium, $c_a^i(\hat f_a) < c_a^i(f_a) \le c_{a'}^i(f_{a'}) \le c_{a'}^i(\hat f_{a'})$ for every arc $a' \in \mathscr{A}_1$ with $\Delta f_{a'} \le 0$, which implies that such an arc cannot be contained in $\hat\sigma(i)$. Therefore, *none* of the arcs in $\mathscr{A}_1$ is contained in $\hat\sigma(i)$. Because also $\hat\sigma(i) \neq e_1 e_3$, the route $\hat\sigma(i)$ must begin with e_2. Because σ and $\hat\sigma$ are equilibria, the inequalities $c_{e_1}^i(f_{e_1}) + c_a^i(f_a) \le c_{e_2}^i(f_{e_2})$ and $c_{e_2}^i(\hat f_{e_2}) \le c_{e_1}^i(\hat f_{e_1}) + c_a^i(\hat f_a)$ hold. However, as the cost functions are strictly increasing and $\Delta f_{e_1}, \Delta f_a > 0 > \Delta f_{e_2}$, these two inequalities contradict each other. A similar contradiction is reached if it is assumed that $\mu(I_{\sigma_2}) < \mu(I_{\hat\sigma_2})$. In this case, there are more users i with $\sigma(i) \in \mathscr{R}_2^-$ and $\hat\sigma(i) = e_2 e_4$ than users i with $\hat\sigma(i) \in \mathscr{R}_2^-$ and $\sigma(i) = e_2 e_4$. Because for every $r \in \mathscr{R}_2^-$, there are more users i with $\hat\sigma(i) \in r$ than users i with $\sigma(i) = r$, this implies that there are some users i with $\hat\sigma(i) \in \mathscr{R}_2^-$ and $\sigma(i) \notin \mathscr{R}_2^-$ for whom $\sigma(i) \neq e_2 e_4$. For each of these users i, $\hat\sigma(i)$ contains some arc $a \in \mathscr{A}_2$ with $\Delta f_a < 0$. Because $\hat\sigma$ is an equilibrium, $c_a^i(f_a) < c_a^i(\hat f_a) \le c_{a'}^i(\hat f_{a'}) \le c_{a'}^i(f_{a'})$ for every arc $a' \in \mathscr{A}_2$ with $\Delta f_{a'} \ge 0$, which implies that such an arc cannot be contained in $\sigma(i)$. Together with the assumption concerning i, this implies that $\sigma(i)$ must begin with e_1. Because σ and $\hat\sigma$ are equilibria, the inequalities $c_{e_1}^i(f_{e_1}) \le c_{e_2}^i(f_{e_2}) + c_a^i(f_a)$ and $c_{e_2}^i(\hat f_{e_2}) + c_a^i(\hat f_a) \le c_{e_1}^i(\hat f_{e_1})$ hold. However, because $\Delta f_{e_1} > 0 > \Delta f_{e_2}, \Delta f_a$, these two inequalities contradict each other. This completes the proof of Claim 2.

Claim 3. $\Delta f_a = 0$ *for all arcs* a.

This is established in Claim 1 under the assumption that $\Delta f_{e_1} \ge 0 \ge \Delta f_{e_3}$. By symmetry, this also holds if $\Delta f_{e_1} \le 0 \le \Delta f_{e_3}$. By Claim 2, Δf_{e_1} and Δf_{e_3} cannot *both* be (strictly) positive. By symmetry, they cannot both be negative, either. Therefore, the assertion of the claim *always* holds. □

Lemma 7.4. *For each of the forbidden networks, there is an assignment of cost functions with two strict Nash equilibria σ and τ such that each user's equilibrium cost in σ is different from that in τ.*

Proof. The cost functions given in Example 1.1 specify such an assignment for the networks in Figures 2(a) and 2(b). The modified version of Example 1.2 given in the paragraph that follows Corollary 4.1 specifies such assignments for the networks in Figures 2(c) and 2(d). □

Lemma 7.5. *Let G' be a network for which there exists an assignment of cost functions with two strict Nash equilibria σ and τ as in Lemma 7.4. Then, the same is true also for every network G'' in which G' is embedded in the wide sense. That is, there is an assignment of cost functions for G'' with two strict Nash equilibria, such that each user's equilibrium cost in one equilibrium is different from that in the other equilibrium.*

Proof. Clearly, it suffices to consider the case in which G'' is obtained from G' by one of the three operations defining embedding in the wide sense. If the operation is the subdivision or addition of an edge, the conclusion is immediate. In the former case, the assignment of cost functions for G'' is the same as for G' except that, for each user, the cost in each direction of the edge that was subdivided is equally divided between its two parts. In the latter case, the cost for each user of the edge that was added is set higher than the user's equilibrium costs in σ and τ. It remains to consider the case in which G'' is obtained from G' by the subdivision of a terminal vertex.

Without loss of generality, the cost of each edge in G' incident with the origin or the destination is greater than unity for all users. Terminal subdivision adds a new edge e, which is incident with the origin or the destination, and is appended to some users' equilibrium routes in σ or in τ. Provided that the cost assigned to e is sufficiently low, the routes with the appended edge remain equilibrium routes for these users, and the two equilibrium costs remain distinct. Specifically, this is the case if the cost of e for each user is less than: (i) unity, (ii) the difference between the cost of the equilibrium route and that of any alternative route in G', both in σ and in τ (these differences are all greater than zero, because the equilibria are strict), and (iii) the absolute value of the difference between the user's equilibrium costs in σ and τ (which, by assumption, are not equal). ☐

Proof of Proposition 4.2. Clearly, for every network G, the first condition implies the second, and the second implies the third. By Lemma 7.1 (or Proposition 2.1), one of the two conditions in Proposition 2.1 holds for G. If (i) in Proposition 2.1 holds, then by Lemmas 7.2 and 7.3, G has the uniqueness property. If (ii) holds, then by Lemmas 7.4 and 7.5, there is an assignment of cost functions for G with two strict Nash equilibria such that each user's equilibrium cost in one equilibrium is different from that in the other. Therefore, either G satisfies all three conditions in Proposition 4.2, or it does not satisfy any of them. ☐

Proof of Theorem 4.1. The proof of the theorem is contained within that of Proposition 4.2. ☐

Proof of Theorem 5.1. Suppose that condition (i) does *not* hold. By the same argument used in the proof of Proposition 4.2, there is an assignment of cost functions for G with two strict Nash equilibria such that the flow on some arc is not the same in both equilibria. Moreover, inspection of the proofs of Lemmas 7.4 and 7.5 shows that there is a partition of the population into three user classes such that this assignment is in the corresponding space $\mathcal{G}$. Because the two equilibria are strict, there is some $\epsilon > 0$ such that each of them is also an equilibrium in every assignment of cost functions in $\mathcal{G}$ that is less than a distance ϵ from the original one. This shows that the set of all assignments of cost functions for which all Nash equilibria are equivalent is *not* dense in $\mathcal{G}$, and thus proves that condition (ii) in the theorem implies (i). It remains to prove the converse implication.

Suppose that the network G satisfies condition (i). Fix some partition of the population, with user classes $I_1, I_2, \ldots, I_n$, and consider the corresponding space of assignments of cost functions $\mathcal{G}$. For each element of $\mathcal{G}$ and each user class I_m, the number of minimal-cost routes for this user class, which will be denoted by φ_m, is the same in all Nash equilibria. This is because the cost for each user of each route is determined by the arc flows, which, by Theorem 4.1, are the same in all Nash equilibria. Therefore, the *mean* number of minimal-cost routes, $\varphi = \sum_{m=1}^{n} \mu(I_m)\varphi_m$, is also the same in all equilibria, and thus defines a real-valued function on $\mathcal{G}$.

Claim 1. *The function* φ: $\mathcal{G} \to \mathbb{R}$ *is upper semicontinuous and has a finite range.*

The proof of the assertion that φ is upper semicontinuous is very similar to that of Milchtaich [12, Lemma 3.4], and is omitted. The second assertion follows from the fact that the cardinality of the range of φ does not exceed $|\mathcal{R}|^n$, the number of routes in G to the power of the number of user classes.

Claim 2. *For every assignment of cost functions in* $\mathcal{G}$ *that is a point of continuity of* φ*, all Nash equilibria are equivalent.*

To prove this claim, consider an assignment of cost functions in $\mathcal{G}$ with two nonequivalent Nash equilibria σ and $\hat{\sigma}$. It has to be shown that φ has a discontinuity at this assignment. For each user class I_m and path p, let f_p^m denote the measure of the set of all users $i \in I_m$

such that p is contained in (or coincides with) $\sigma(i)$, and $\hat{f}_p^m$ the corresponding quantity for $\hat{\sigma}$. Because the two equilibria are not equivalent, there is a user class I_{m_0} such that $f_{a_0}^{m_0} \neq \hat{f}_{a_0}^{m_0}$ for some arc a_0. By assumption, G is nearly parallel or consists of several nearly parallel networks connected in series. The arc a_0 is in one of these networks, G'. Without loss of generality, it may be assumed that, for each user i, the routes $\sigma(i)$ and $\hat{\sigma}(i)$ coincide outside G'. (If this is not so, $\hat{\sigma}$ can be replaced by another Nash equilibrium, in which the users' routes agree with their routes in $\hat{\sigma}$ inside G' and with those in σ outside G'.) Let $\mathcal{R}'$ denote the set of all routes in G'. The arc a_0 is contained in some $r_0 \in \mathcal{R}'$ with $f_{r_0}^{m_0} \neq \hat{f}_{r_0}^{m_0}$. Therefore, there is a real number α such that the affine combination $\tilde{f}_r^m \stackrel{\text{def}}{=} \alpha f_r^m + (1-\alpha)\hat{f}_r^m$ is nonnegative for all user classes I_m and $r \in \mathcal{R}'$, and zero for some I_{m_1} and $r_1 \in \mathcal{R}'$ with $f_{r_1}^{m_1} \neq \hat{f}_{r_1}^{m_1}$. Because the population measure μ is nonatomic and $\sum_{r\in\mathcal{R}'} \tilde{f}_r^m = \mu(I_m)$ for all user classes I_m, there is a strategy profile $\tilde{\sigma}$ that coincides with $\sigma(i)$ and $\hat{\sigma}(i)$ outside G' and has the property that, for every user class I_m and $r \in \mathcal{R}'$, the measure of the set of all users $i \in I_m$ such that r is contained in (or coincides with) $\tilde{\sigma}(i)$ is equal to $\tilde{f}_r^m$. Because by Theorem 4.1, $\sum_{m=1}^n f_a^m = \sum_{m=1}^n \hat{f}_a^m$ for all arcs a, the flow on each arc in $\tilde{\sigma}$ is the same as in σ and $\hat{\sigma}$. For every user class I_m and $r \in \mathcal{R}'$, $\tilde{f}_r^m > 0$ only if $f_r^m > 0$ or $\hat{f}_r^m > 0$, and hence, only if r is contained in some minimal-cost route for user class I_m. Therefore, $\tilde{\sigma}$ is a Nash equilibrium. Before the proof of Claim 2 can be completed, the following has to be established.

Claim 3. *There is some arc a_1 in G' that is contained in r_1 and for which $f_{a_1}^{m_1} > 0$ or $\hat{f}_{a_1}^{m_1} > 0$, but $\tilde{f}_{a_1}^{m_1} = 0$.*

This is easily shown if G' is homeomorphic to one of the networks in Figures 4(a)–4(d). In each of these networks, each route contains some arc that is not contained in any other route. Because by construction $f_{r_1}^{m_1} \neq \hat{f}_{r_1}^{m_1}$ and $\tilde{f}_{r_1}^{m_1} = 0$, this implies that similar inequality and equality hold with the route r_1 in G' replaced by one of its arcs a_1. If G' is homeomorphic to the network in Figure 4(e), there are two cases to consider. If the two nonterminal vertices with degree three or higher, u and v, are both included in r_1, then r_1 contains an arc a_1 that is not contained in any other route in G'. Therefore, this case is similar to the one considered above. It remains to consider the case in which only u, say, is included in r_1. In this case, there is a unique route r_2 in G' that does not share any arc with r_1, as well as at least one route r_3 that shares with r_1 only the arcs that follow u, and at least one route r_4 that shares with r_1 only the arcs that precede u. Any such pair of routes r_3 and r_4 together contain all the arcs contained in r_1 and r_2, and at least two additional arcs. Therefore, if $\tilde{f}_{r_3}^{m_1} > 0$ and $\tilde{f}_{r_4}^{m_1} > 0$, the sum of the costs of r_3 and r_4 for user class I_{m_1} exceeds that of r_1 and r_2, which implies that at least one of them is not contained in any minimal-cost route. However, it is shown above that $\tilde{f}_r^m > 0$ (for $r \in \mathcal{R}'$) implies that r *is* contained in some minimal-cost route for user class I_m. Therefore, $\tilde{f}_{r_3}^{m_1} = 0$ or $\tilde{f}_{r_4}^{m_1} = 0$. It follows, as also $\tilde{f}_{r_1}^{m_1} = 0$, that r_1 contains at least one arc a_1 with $\tilde{f}_{a_1}^{m_1} = 0$. Because by assumption $f_{r_1}^{m_1}$ is not equal to $\hat{f}_{r_1}^{m_1}$, at least one of them is not zero. If $f_{r_1}^{m_1} > 0$, then $f_{a_1}^{m_1} > 0$, and if $\hat{f}_{r_1}^{m_1} > 0$, then $\hat{f}_{a_1}^{m_1} > 0$. This completes the proof of Claim 3.

The proof of Claim 2 can now be completed. Any arc a_1 as in Claim 3 is included in at least one (since $f_{a_1}^{m_1} > 0$ or $\hat{f}_{a_1}^{m_1} > 0$), but not all (since $\tilde{f}_{a_1}^{m_1} = 0$), minimal-cost routes for user class I_{m_1}. For any $\epsilon > 0$, $\tilde{\sigma}$ is a Nash equilibrium also with respect to the assignment of cost functions obtained from that considered above by adding ϵ to each cost function $c_a^i(x)$ with $i \in I_{m_1}$ and $a = a_1$, and leaving all the other cost functions unchanged. The set of minimal-cost routes for user class I_{m_1} in the new assignment is a proper subset of that in the original assignment. It consists of all the minimal-cost routes for the original assignment that do not contain a_1. For any other user class, the two sets are equal. Therefore, the value of φ for the new assignment is smaller than the original value by at least $\mu(I_{m_1})$. Because the distance between the two assignments is ϵ, and hence can be chosen to be arbitrarily small, this proves that φ has a discontinuity at the original assignment. This completes the proof of Claim 2.

Together with Claims 1 and 2, the following claim completes the proof of the theorem.

CLAIM 4. *In every metric space $\mathscr{X}$, the set of all points of continuity of an upper semicontinuous function g: $\mathscr{X} \to \mathbb{R}$ with a finite range is open and dense.*

Every point of continuity of such a function g has an open neighborhood in which this function is constant, and hence continuous. This proves that the set of all points of continuity is open. To prove that this set is dense in $\mathscr{X}$, consider any open set U. Let $x_0 \in U$ be such that $g(x_0) = \min_{x \in U} g(x)$. By upper semicontinuity of g, there is a neighborhood V of x_0 such that $g(x) \leq g(x_0)$ for all $x \in V$. Clearly, in $U \cap V$, $g(x) = g(x_0)$ for all x, and g is therefore continuous at x_0. $\square$

Acknowledgments. The author thanks Hideo Konishi, Michael Ben-Or, and two anonymous referees for helpful comments and suggestions. This research was supported by the Israel Science Foundation (Grant 749/02).

References

[1] Aashtiani, H. Z., T. L. Magnanti. 1981. Equilibria on a congested transportation network. *SIAM J. Algebraic Discrete Methods* **2** 213–226.
[2] Altman, E., H. Kameda. 2001. Equilibria for multiclass routing in multi-agent networks. *Proc. 40th IEEE Conf. on Decision and Control*, Orlando, FL, 604–609.
[3] Beckmann, M., C. B. McGuire, C. B. Winsten. 1956. *Studies in the Economics of Transportation*. Yale University Press, New Haven, CT.
[4] Bell, M .G. H., Y. Iida. 1997. *Transportation Network Analysis*. Wiley, Chichester, UK.
[5] Dafermos, S. C. 1972. The traffic assignment problem for multiclass-user transportation networks. *Transportation Sci.* **6** 73–87.
[6] Diestel, R. 2000. *Graph Theory*, 2nd ed. *Graduate Texts in Mathematics*, Vol. 173. Springer-Verlag, New York.
[7] Harary, F. 1969. *Graph Theory*. Addison-Wesley, Reading, MA.
[8] Konishi, H. 2004. Uniqueness of user equilibrium in transportation networks with heterogeneous commuters. *Transportation Sci.* **38** 315–330.
[9] Mas-Colell, A. 1985. *The Theory of General Economic Equilibrium: A Differentiable Approach*. Cambridge University Press, Cambridge, UK.
[10] Milchtaich, I. 1996. Congestion games with player-specific payoff functions. *Games Econom. Behavior* **13** 111–124.
[11] Milchtaich, I. 1998. Crowding games are sequentially solvable. *Internat. J. Game Theory* **27** 501–509.
[12] Milchtaich, I. 2000. Generic uniqueness of equilibrium in large crowding games. *Math. Oper. Res.* **25** 349–364.
[13] Milchtaich, I. 2005. Network topology and the efficiency of equilibrium. Mimeo.
[14] Nagurney, A. 1999. *Network Economics: A Variational Inequality Approach*, 2nd ed. Kluwer, Boston, MA.
[15] Newell, G. F. 1980. *Traffic Flow on Transportation Networks*. MIT Press, Cambridge, MA.
[16] Rath, K. P. 1992. A direct proof of the existence of pure strategy equilibria in games with a continuum of players. *Econom. Theory* **2** 427–433.
[17] Schmeidler, D. 1970. Equilibrium points of nonatomic games. *J. Statist. Phys.* **7** 295–300.
[18] Sheffi, Y. 1985. *Urban Transportation Networks*. Prentice-Hall, Englewood Cliffs, NJ.
[19] Wardrop, J. G. 1952. Some theoretical aspects of road traffic research. *Proc. Inst. Civil Engineers* **1**(Part II) 325–378.