

Evolutionary Implementation and Congestion Pricing

WILLIAM H. SANDHOLM
University of Wisconsin

First version received 4 January 2000; final version accepted 4 December 2001 (Eds.)

We consider an implementation problem faced by a planner who manages a roadway network. The problem entails both hidden information and hidden actions. We solve the planner's problem by introducing a new class of mechanisms and a new notion of implementation. The mechanisms, called price schemes, attach transfers to the available routes; they do not involve direct revelation. The method of implementation is evolutionary, requiring that players who follow any reasonable myopic adjustment process eventually learn to behave as the planner desires. We show that efficient behaviour can be guaranteed using simple, decentralized price schemes.

1. INTRODUCTION

Roadway congestion is a source of enormous economic costs.¹ In principle, many of these costs can be prevented, as they result from socially inefficient choices by individual drivers. For example, it is socially optimal for commuters who only receive small benefits from driving to work not to drive at all, as these benefits are outweighed by the delays they create for others. However, since these drivers ignore the externalities they create, they may opt to drive anyway. It is also socially desirable to limit traffic on narrow, easily congested roads: the externalities created by drivers on these roads are more severe than those which develop on wider expressways. But drivers typically will choose to drive on narrow roads if they offer shorter travel times. Such behaviour creates needlessly high congestion levels.

A number of regions have considered alleviating roadway congestion by introducing congestion pricing. The pioneer in this regard is Singapore, which introduced the first congestion pricing scheme in 1975.² Originally, this scheme set a toll for entering the central city during peak driving hours; the scheme was administered using window stickers which were checked by police at certain checkpoints. In 1998, Singapore began to administer its scheme using an electronic collection system. Under this system, gantries are stationed over roadways at the points where tolls are charged. The gantries send signals to receivers installed within individual vehicles; tolls are deducted from the driver's debit card. This system is designed to generate an error rate of less than 1 in 100,000 even when traffic is heavy and moving at speeds of 120 km h^{-1} , so that the toll collection process itself is not a source of delay.³ Moreover, electronic tolling makes it possible

1. Schrank and Lomax (1999) estimate that in 1997, roadway congestion in the ten most congested U.S. cities wasted 2.6 billion hours of drivers' time, yielding a dollar cost of \$39.3 billion. In Los Angeles alone, the corresponding figures were 739 million hours and \$10.8 billion, respectively.

2. More recently, France and California have begun to subject certain specific routes to congestion pricing. The Norwegian cities of Bergen, Oslo and Trondheim possess toll rings around their city centres; while originally introduced to finance infrastructure improvements, congestion management has become a secondary goal. Many cities and regions have undertaken comprehensive studies of the impact of congestion pricing, most notably Hong Kong, Cambridge (U.K.), Stockholm, the Randstad (Holland), and London. For further details on plans for and introductions of congestion pricing schemes, see Gomez-Ibañez and Small (1994), *The Economist* (1997), and Small and Gomez-Ibañez (1998).

3. For more on congestion pricing in Singapore and on electronic tolling technologies, see Gomez-Ibañez and Small (1994), Phang and Toh (1997) and Soo (1998).

to place charges on specific road segments, and for charges to be varied with current congestion levels. These features bring substantial efficiency gains within reach.

In this paper, we view the alleviation of network congestion as an implementation problem. As is usual in such problems, we suppose that the planner must contend with hidden information: he does not know individual drivers' tolerances for delay. We also suppose that the planner faces a hidden action problem: he only has a limited ability to observe drivers' behaviour. As we shall see, the hidden actions can render dominant strategy implementation impossible in this setting.

We solve the planner's problem by introducing a new class of mechanisms and a new notion of implementation. The mechanisms, called price schemes, attach transfers to the available routes; they do not involve direct revelation. The notion of implementation is evolutionary, requiring that players who follow any reasonable myopic adjustment process learn to behave efficiently. We show that there are simple, decentralized price schemes which ensure efficient play while respecting the planner's limited knowledge of players' preferences and behaviour.

We can describe the planner's problem in the following way. A collection of towns is connected by a network of streets. Commuters in these towns must travel from their homes to their offices. The time such a journey requires is the sum of the delays on each street the driver takes; the delay on each street depends on the number of drivers on that street. Different drivers will tolerate different levels of delay before preferring to stay home. The planner would like to ensure that the right drivers decide to commute, and that those who do commute use the network in the most efficient manner possible.

We assume that the planner is faced with two constraints. First, he does not know players' valuations for completing their commutes. Were this the planner's only constraint, he could solve his problem by applying a standard revelation mechanism, namely the Vickrey–Clarke–Groves (VCG) mechanism (Vickrey (1961), Clarke (1971), Groves (1973)). Under this mechanism, each player reports a valuation to the planner, who for each profile of reports specifies an allocation of drivers over routes and the transfers paid by each player. The allocation is chosen to be efficient conditional on reports being truthful, while the transfers are chosen to make truth-telling a dominant strategy. Thus, were hidden information is the planner's only constraint, he could implement efficient behaviour in dominant strategies.

In employing the VCG mechanism, one implicitly assumes that the distribution of drivers over routes specified by the mechanism will be followed by the players. This assumption might be justified if the planner were able to observe each driver's choice, as compliance could then be guaranteed using a forcing contract. However, when the number of drivers is large, it seems unrealistic to assume that the planner can perfectly observe their choices. If he cannot, the players may prefer to ignore his recommendation. For example, a player told to take a roundabout route to work may instead take a shorter route, confounding the planner's attempt to ensure efficient behaviour.

Rather than assume that the planner has perfect knowledge of behaviour, it seems more reasonable to suppose that behaviour is *anonymous*: only aggregate behaviour can be observed, and transfers can depend on this and on action (*i.e.* route) choices, not on players' names.⁴ To study implementation when behaviour is anonymous, we must include the players' route choices explicitly in the mechanism design problem, and must restrict the planner to mechanisms which respect the players' anonymity. The simplest such mechanisms are *price schemes*. Under these mechanisms, each player chooses a route to work or chooses to stay home; the planner attaches

4. The hidden action problem described here, which is due to a coarse information structure, differs from the usual one in information economics, in which the principal receives a noisy signal concerning the agents' behaviour. Myerson (1982) proves a (Bayesian equilibrium) revelation principle for mechanism design problems with both hidden information and hidden actions. His paper focuses on incentive compatibility, and does not address the problem of multiple equilibria.

prices to each of these actions which may depend on the players' aggregate behaviour, but which may not depend on the name of the player making the choice.

Unfortunately, when behaviour is anonymous, dominant strategy implementation may be impossible. To see this most easily, suppose that all players attach the same value to getting to work. Also assume the planner chooses to employ a price scheme.⁵ Suppose that under this scheme a certain driver judges a particular route to be a dominant strategy: regardless of what the others do, the driver prefers to drive on this route. Absent idiosyncratic preferences between the streets themselves, it follows that all players who opt to drive will choose this same route. Therefore, any social choice function which sometimes requires different players to take different routes cannot be implemented in dominant strategies.

In settings where dominant strategy implementation is known to be impossible, implementation using Nash equilibrium or one of its refinements is often considered next. This approach implicitly assumes that players will follow an equilibrium of the mechanism which the planner provides. In order to justify the equilibrium assumption via a rationalistic approach, one needs players with considerable knowledge of their opponents' intentions and reasoning abilities.⁶ However, when there are large numbers of players, as is the case in the network planner's problem, these assumptions seem especially strong, and so predictions which rely on them may not have much force.

In order to find a middle ground between dominant strategy and Nash implementation, we consider a notion of implementation based on evolutionary game theory. In particular, we suppose that once a price scheme is in place, the players adjust their behaviour over time in response to current delays and to the incentives which the price scheme provides. The adjustment process itself is described by a differential equation. Rather than specify a particular functional form for this equation (e.g. the replicator dynamic), we instead ask that it be a member of a broad class of admissible dynamics. The main restriction defining this class is quite weak, requiring only that strategies' growth rates be positively correlated with their payoffs.

The planner would like to ensure that regardless of their types, the players learn to follow the efficient driving pattern. A price scheme *globally implements* the efficient social choice function if for every possible type profile, the efficient driving pattern is globally stable under all admissible dynamics.⁷

We show that the planner can solve his implementation problem by selecting a mechanism from a class whose members we call *variable price schemes*.⁸ The schemes we consider are *separable*, in the sense that the price of any route can be decomposed into prices on each individual street in the route, where the price of each street only depends on the number of drivers on that street. Separability implies that the variable price schemes can be imposed in a decentralized fashion.

The collection of variable price schemes is indexed by a parameter called the *elasticity threshold*, which specifies the lowest level of sensitivity to congestion which a street must exhibit to be assigned a positive price. When demand to use the network is inelastic, the elasticity threshold determines the revenues generated by the tolling scheme. When demand is elastic,

5. An analogous argument can be applied to other mechanisms which respect the players' anonymity.

6. See, for example, Aumann and Brandenburger (1995).

7. As we shall see, global implementation implies implementation in Nash equilibrium. However, the former notion of implementation is considerably stronger, as it requires not only that there be a unique Nash equilibrium outcome, but also that this outcome be globally stable under any reasonable behaviour adjustment process.

8. The word "variable" refers to the fact that the prices charged are not set in advance, but rather are functions of current network utilization.

different thresholds generate efficiency under different welfare measures, which vary the relative weights assigned to costs and benefits in assessing a driver allocation.⁹

Variable price schemes can be viewed as generalizations of marginal cost pricing. “Marginal cost pricing” usually refers to an equilibrium phenomenon: that by making agents pay for the externalities they create in equilibrium, one can guarantee the efficiency of equilibrium play. In contrast, variable price schemes set prices appropriately both in and out of equilibrium. By doing so, a planner can render efficient behaviour the unique equilibrium and the global attractor of any reasonable adjustment process. More importantly, he can do so without knowing anything about players’ types, and hence without being able to predict what the equilibrium will turn out to be.

By imposing a variable price scheme and noting where behaviour settles, the planner can learn the efficient allocation of drivers over roads. After determining this allocation, the planner can maintain efficient behaviour using a *fixed price scheme*, which *always* charges the tolls imposed under some variable price scheme at the efficient allocation. This simple scheme ensures that the efficient allocation is globally stable, so that play will return to this allocation after any shock to behaviour.

Our analysis is based on results on potential games developed in Sandholm (2001). A potential game is a game which admits a potential function: a real valued function defined on the space of strategy distributions which any reasonable evolutionary process must ascend.¹⁰ By executing a variable price scheme, the social planner ensures that regardless of the realization of demand, the players face a potential game whose potential function is proportional to the social welfare measure. Because of this, we can establish that under the variable price schemes, players always learn to make socially optimal decisions about both whether and how to play the game.

Recent papers by Cabrales (1999), Cabrales and Ponti (2000) and Ponti (2000) use techniques from evolutionary game theory to evaluate well-known mechanisms from the implementation literature, obtaining mixed results.¹¹ Cabrales (1999) shows that a modified version of the Nash implementation mechanism of Repullo (1987) has strong evolutionary stability properties. He also establishes an instability result for Abreu and Matsushima’s (1994) mechanism, which relies on the iterated elimination of weakly dominated strategies. Similarly, Cabrales and Ponti (2000) show that Sjöström’s (1994) mechanism can possess weakly dominated Nash equilibria which do not yield the desired outcome but which are limit points of payoff monotone evolutionary dynamics. Ponti (2000) proves a related result for Glazer and Ma’s (1989) solution to King Solomon’s dilemma, and provides an alternative mechanism with better evolutionary properties.

The papers described above all concern implementation in settings with hidden information but without hidden actions, in which well-known mechanisms can be employed. In contrast, to address the network planner’s problem we must devise new mechanisms which explicitly

9. If equal weights are chosen, the welfare measure is simply consumer surplus. However, other weights are sometimes desirable. For example, driving causes pollution, an externality which is not directly accounted for in our model. If the planner values clean air, he may want to attach extra importance to lowering aggregate driving time when assessing the welfare of a driving pattern.

10. An early application of potential functions can be found in Beckmann *et al.* (1956), who use a potential function to characterize equilibria in an elastic demand traffic model. Rosenthal (1973) introduces congestion games with finite numbers of players and uses a potential function argument to establish the existence of a pure strategy equilibrium. Building on the latter paper, Monderer and Shapley (1996) define finite player potential games. They show that maximizers of potential are both the only equilibria of these games and the only possible limits of better reply strategy adjustment processes. In Sandholm (2001), we define infinite player potential games and prove related results characterizing equilibrium and evolution; some of these results are presented below. We also establish conditions under which equilibria are efficient, and characterize our infinite player potential games as the limits of Monderer and Shapley’s (1996) finite player games.

11. Other papers which address the dynamics of implementation include Muench and Walker (1983), Walker (1984), Jordan (1986), de Trenchalaye (1988, 1989) and Vega-Redondo (1989); see Cabrales (1999) for a discussion.

allow for the players' anonymous choices. We exhibit pricing mechanisms which possess strong evolutionary stability properties, and which seem possible to administer using current tolling methods.

We open our analysis by studying price schemes for network congestion games with inelastic demand. Section 2 introduces congestion games, potential games, and admissible evolutionary dynamics. Section 3 describes the planner's problem and defines our evolutionary notion of implementation. Section 4 presents the variable and fixed price schemes, showing that the former can be used to ensure efficient play when demand to use the network is unknown, and that the latter can be used to maintain efficient play after the efficient allocation is discovered. Sections 5–7 present definitions and results for the more complicated case of elastic demand. Section 8 concludes.

2. GAMES AND EVOLUTION UNDER INELASTIC DEMAND

2.1. Games with continuous player sets

Each player in a game with continuous player sets is a member of a population in the set $R = \{1, \dots, \bar{r}\}$, each element of which corresponds to a different player role. A typical population is denoted $r \in R$. The mass of each population is m^r , and the vector $m = (m^1, \dots, m^{\bar{r}})$ lists the masses of all populations.

We let S^r be the set of strategies available to population r , and let n^r be the number of such strategies. We also let $S = \bigcup_{r \in R} S^r$ denote the *union* of all populations' strategy sets, and let $n = \sum_{r \in R} n^r$ equal the total number of strategies available to all populations.

A typical *strategy distribution* is $x \in X_m = \{x \in \mathbf{R}_+^n : \sum_{i \in S^r} x_i = m^r \text{ for all } r\}$, where x_i is the mass of players in population r who choose strategy $i \in S^r$. The payoff function for strategy i is denoted $F_i: X_m \rightarrow \mathbf{R}$, while the payoff functions for all strategies are collectively denoted $F: X_m \rightarrow \mathbf{R}^n$. A strategy distribution $x \in X_m$ is a *Nash equilibrium* if all players choose strategies which maximize their payoffs given their opponents' behaviour:

$$F_i(x) = \max_{j \in S^r} F_j(x) \quad \text{whenever } i \in S^r \quad \text{and} \quad x_i > 0.$$

2.2. Congestion games

Consider a group of drivers who live in a collection of towns connected by a network of streets. Each driver must commute from his hometown to the town of his workplace. He does so by selecting a route (*i.e.* a subset of the streets) leading from home to work. A driver's total travel time is the sum of the delays on each street, which are each increasing functions of the number of drivers on that street.

An inelastic demand congestion model is a collection $\{R, \{m^r\}_{r \in R}, \{S^r\}_{r \in R}, \{\Phi_i\}_{i \in S}, \{c_\phi\}_{\phi \in \Phi}\}$. R is a set of one or more populations, one for each home/work location pair. The finite set $\Phi = \bigcup_{i \in S} \Phi_i$ contains all available streets. Each strategy $i \in S^r$ corresponds to a complete route (*i.e.* a non-empty collection of streets) $\Phi_i \subset \Phi$ which connects the home and work pair r .¹²

Let $\rho(\phi) = \{i \in S : \phi \in \Phi_i\}$ denote the set of routes which require street ϕ . The *utilization* of street $\phi \in \Phi$ is the total mass of the players who drive on that street:

$$u_\phi(x) = \sum_{i \in \rho(\phi)} x_i.$$

12. We do not assume any graph theoretic structure on the set of streets Φ . Hence, all of our results can be used to study congestion in settings in which the facilities $\phi \in \Phi$ are not arranged in a network.

The cost functions $c_\phi : \mathbf{R}_+ \rightarrow \mathbf{R}$ report the delay on a street as a function of the number of drivers using that street. As some streets (e.g. the narrower ones) are more prone to delays than others, different streets will typically have different cost functions. We assume that each cost function c_ϕ is non-negative, continuously differentiable, and strictly increasing: $c_\phi(u) \geq 0$ and $c'_\phi(u) > 0$ for all u .¹³

The total delay from using route $i \in S$ is determined by adding the delays on the streets in the route. Hence, $C_i: X_m \rightarrow \mathbf{R}$, the cost of using strategy i as a function of the overall strategy distribution, is given by

$$C_i(x) = \sum_{\phi \in \Phi_i} c_\phi(u_\phi(x)).$$

To define a *congestion game*, we must specify payoff functions for all strategies. Here, payoffs are simply the negations of delay costs: $F_i(x) = -C_i(x)$.

To describe the inelastic demand implementation problem, we suppose that at some initial stage, Nature specifies that each player is of one of two types: “stay home” or “commute”. The demand vector $m = (m^1, \dots, m^{\bar{r}})$ represents the realized numbers of active commuters in each population.

Because the planner does not know which realization of m occurs, payoffs must be defined for all possible realizations. If the number of potential commuters in population r is M^r , then the realized number of active commuters will lie between 0 and M^r . Hence, the set of possible demand vectors is $\mathbf{M} = \prod_{r \in R} [0, M^r]$, while the set of possible strategy distributions is $X = \bigcup_{m \in \mathbf{M}} X_m = \{x \in \mathbf{R}_+^n : \sum_{i \in S^r} x_i \leq M^r \text{ for all } r\}$. Fortunately, our earlier definitions of the cost and payoff functions on the set X_m do not depend on the demand vector m ; therefore, these definitions extend immediately to all of X .¹⁴

For convenience, we assume throughout the paper that *strategy distributions are distinguishable*: for each distinct pair $x, y \in X$, there is a street $\phi \in \Phi$ such that $u_\phi(x) \neq u_\phi(y)$. While this assumption simplifies our analysis, it is not essential—see footnote 18 below.

2.3. Evolutionary dynamics

The rationalistic approach to justifying Nash equilibrium requires players to know their opponents’ intentions. When the number of players is large this requirement is quite strong. But if the game is played repeatedly, we can avoid this assumption by modelling behaviour as a myopic adjustment process during which players switch to strategies which improve their current payoffs.

An evolutionary dynamic is described by a vector field $V: X \rightarrow \mathbf{R}^n$. This vector field defines an equation of motion $\dot{x} = V(x)$ on the space of strategy distributions. We call V *admissible* with respect to the game F if it satisfies the following five conditions:

- V is Lipschitz continuous. (LC)
- $V_i(x) \geq 0$ whenever $x_i = 0$. (FI 1)
- $\sum_{i \in S^r} V_i(x) = 0$ for all $x \in X$ and $r \in R$. (FI 2)
- $V(x) \cdot F(x) > 0$ whenever $V(x) \neq \vec{0}$. (PC)
- $V(x) = \vec{0}$ implies that x is an equilibrium of F . (NC)

13. Allowing c'_ϕ to equal zero has only a minor impact on our analysis.

14. When they are defined on all of X , the payoffs F do not define a single congestion game, but a collection of congestion games, one for each demand vector m . Nevertheless, we will sometimes abuse terminology and refer to F itself as a game.

The first three conditions are technical requirements which ensure the existence of unique solution trajectories which stay in the space X .¹⁵ The more important conditions are the last two, which link the dynamic to the game’s payoffs. To interpret condition (PC), note that by forward invariance condition (FI 2),

$$\begin{aligned} V(x) \cdot F(x) &= \sum_{r \in R} \sum_{i \in S^r} V_i(x) F_i(x) \\ &= \sum_{r \in R} \left(\sum_{i \in S^r} (V_i(x) - 0) \left(F_i(x) - \frac{1}{n^r} \sum_{j \in S^r} F_j(x) \right) \right) \\ &= \sum_{r \in R} n^r \text{Cov}(V^r(x), F^r(x)), \end{aligned}$$

where $\text{Cov}(V^r, F^r)$ is the covariance of growth rates and payoffs in population r . Thus, condition (PC) requires a *positive correlation* between growth and payoffs. While this condition holds if there is a positive covariance in each population, it only requires the weighted sum of the covariances of all populations to be positive. Positive correlation is the weakest monotonicity condition used in the evolutionary game theory literature; for some comparisons, see Sandholm (2001).

We call condition (NC) *noncomplacency*. It requires that a population that is not playing a Nash equilibrium continues to adjust its behaviour. When behaviour is not in equilibrium, there are players who would benefit from switching strategies; noncomplacency requires that some players eventually avail themselves of this opportunity.¹⁶

2.4. *Potential games*

We now review results from Sandholm (2001) which we need to analyze our price schemes. We call a game with continuous player sets a *potential game* if there is a function f which satisfies

$$\frac{\partial f}{\partial x_i}(x) = F_i(x) \quad \text{for all } x \in X \quad \text{and} \quad i \in S.$$

We call f the game’s *potential function*. If a game admits a potential function, this function can be used to determine the game’s Nash equilibria and to characterize evolutionary dynamics. In particular, Sandholm (2001), building on the work of Beckmann *et al.* (1956), Rosenthal (1973) and Monderer and Shapley (1996), establishes the following result.

Lemma 1. (i) *Let F be a potential game, and let V be a dynamic which is admissible with respect to F . Then every solution trajectory of V converges to a connected set of Nash equilibria of F .*

(ii) *Suppose that in addition, the potential function f of F is strictly concave, and fix a demand vector m . Then the maximizer of f on the set X_m is the unique Nash equilibrium of F in X_m and is the global attractor under V of all trajectories in X_m .*

For completeness, the proof of this result is presented in the appendix.

In general, evolutionary dynamics of games need not converge to Nash equilibria: solution trajectories may converge to non-Nash rest points or to limit cycles, and can even exhibit

15. Condition (FI 1) requires that the mass of players using each strategy never becomes negative, and condition (FI 2) requires that the mass of players in each population remains constant.

16. An example of an admissible dynamic is the Brown–von Neumann–Nash (BNN) dynamic; see Brown and von Neumann (1950), Weibull (1996), Berger and Hofbauer (2000) or Sandholm (2001). The replicator dynamic is not admissible because it fails condition (NC). However, our analysis can be extended to this dynamic—see the discussion at the end of the next section.

chaotic behaviour. However, if a game admits a potential function, all solution trajectories of all admissible dynamics converge to connected sets of Nash equilibria.¹⁷ If the potential function is strictly concave, the equilibrium is unique, and therefore globally stable.¹⁸

For intuition, observe that the definitions of potential and positive correlation imply that any solution trajectory $\{x_t\}_{t \geq 0}$ of an admissible dynamic V must satisfy

$$\frac{d}{dt} f(x_t) = \nabla f(x_t) \cdot \dot{x}_t = F(x_t) \cdot V(x_t) \geq 0.$$

That is, profitable behaviour adjustments must increase potential. Local maximizers of potential reflect the absence of profitable adjustments, and hence are Nash equilibria. If the potential function is strictly concave, profitable adjustments must lead to its unique maximizer, which is the unique equilibrium of the game.

It is easily verified that all congestion games admit the potential function

$$f(x) = - \sum_{\phi \in \Phi} \int_0^{u_\phi(x)} c_\phi(z) dz.$$

Lemma 2 shows that when costs are strictly increasing, this function is strictly concave. Hence, Lemma 1 implies that in any congestion game in which congestion is a “bad”, there is a unique, globally stable equilibrium for each demand vector.

Lemma 2. *Any congestion game with distinguishable strategy distributions and with cost functions satisfying $c'_\phi > 0$ has a strictly concave potential function.*

Proof. Let $x, y \in X$; it is enough to show that the potential function is strictly concave on the line segment joining x and y . For $\lambda \in [0, 1]$, let $z(\lambda) = \lambda x + (1 - \lambda)y$, and let $z_\phi(\lambda) = u_\phi(\lambda x + (1 - \lambda)y)$. Then since $u_{\hat{\phi}}(x) \neq u_{\hat{\phi}}(y)$ for some $\hat{\phi} \in \Phi$ by distinguishability, $z'_{\hat{\phi}}(\lambda) = u_{\hat{\phi}}(x) - u_{\hat{\phi}}(y) \neq 0$ for this $\hat{\phi}$, and so

$$\frac{d^2}{(d\lambda)^2} f(z(\lambda)) = - \sum_{\phi \in \Phi} c'_\phi(z_\phi(\lambda)) (z'_\phi(\lambda))^2 < 0. \quad \parallel$$

In restricting attention to dynamics which satisfy the noncomplacency condition (NC), we rule out the replicator dynamic. This dynamic can be interpreted as a model of evolution through imitation of successful agents.¹⁹ Consequently, strategies which are initially absent from the population are never used, and so the dynamic admits non-Nash rest points on the boundary of the state space. In the appendix, we consider evolution under *payoff monotone* dynamics, a class of dynamics which includes the replicator dynamic but whose members all violate condition (NC). We show (Proposition A1) that in games with a strictly concave potential function, all solution trajectories of these dynamics from interior initial conditions converge to the game’s unique Nash equilibrium. Hence, versions of all of the results which follow can be proved for payoff monotone dynamics under this restriction on initial behaviour.

17. It is known that if an interior solution trajectory of a payoff monotone dynamic (see below) converges to a unique limit point, this point must be a Nash equilibrium, regardless of whether the underlying game is a potential game. In potential games, setwise convergence of all solution trajectories of all admissible dynamics is guaranteed.

18. If the potential function is concave but not strictly so, equilibrium need not be unique, but the set of Nash equilibria is equal to the closed, convex set of maximizers of potential. For this reason, the strict concavity assumptions used below are not essential. In fact, the only role of our distinguishability assumption is to ensure that concavity of potential holds strictly (see the proof of Lemma 2), so this assumption is dispensable as well.

19. See Björnerstedt and Weibull (1996) and Schlag (1998).

3. EVOLUTIONARY IMPLEMENTATION UNDER INELASTIC DEMAND

Equilibria of congestion games are typically inefficient, as easily congested streets tend to be overused relative to the social optimum. The planner would like to ensure that players learn to behave efficiently, but must do so without knowing demand. Furthermore, because the players' behaviour is anonymous, the planner cannot simply tell each player how to act, but must influence their choices through some incentive scheme which respects their anonymity.

One way the planner can influence behaviour while respecting these restrictions is to set prices for the use of each route. Formally, a *price scheme* is a function $P: X \rightarrow \mathbf{R}^n$. Its components $P_i: X \rightarrow \mathbf{R}$ represent the payments a player choosing route i must make to the planner as a function of society's aggregate behaviour. By introducing a price scheme, the planner creates a new game whose payoffs are

$$\hat{F}_i(x) = F_i(x) - P_i(x) = -C_i(x) - P_i(x).$$

The planner would like to choose prices in such a way that regardless of the demand vector m , myopic adjustment with respect to the new payoffs \hat{F} leads to efficient behaviour with respect to the original payoffs F .

To express this idea more precisely, we define a *social choice function* $\sigma: \mathbf{M} \rightarrow X$ to be a map which specifies a strategy distribution $\sigma(m) \in X_m$ for each demand vector m . The price scheme P *globally implements* the social choice function σ if for each demand vector m , the distribution $\sigma(m)$ is globally stable under any dynamic V which is admissible with respect to the augmented game \hat{F} . By introducing a price scheme which globally implements the efficient state, the planner ensures that regardless of its demand or its initial behaviour, a population of players who revise their choices in a reasonable fashion will learn to behave efficiently.

If the price scheme P globally implements the social choice function σ , then for each demand vector m , the strategy distribution $\sigma(m)$ is the unique rest point in X_m of any admissible dynamic, and is hence the game's unique Nash equilibrium (see Lemma A1 in the appendix). Therefore, global implementation implies Nash implementation. However, since even a unique Nash equilibrium may not be globally or locally stable under admissible dynamics, the converse implication is false: of the two notions of implementation, global implementation is strictly more demanding.

In principle, the price $P_i(x)$ for route i may depend on the entire strategy distribution x without violating the players' anonymity. But since x lists the numbers of players choosing each complete route, keeping track of the full strategy distribution is a demanding task. We therefore require prices to be of the form

$$P_i(x) = \sum_{\phi \in \Phi_i} p_\phi(u_\phi(x))$$

for some functions $p_\phi: \mathbf{R}_+ \rightarrow \mathbf{R}$. We call a price scheme *separable* if it can be decomposed in this way. Under a separable scheme, the total price P_i of each route i can be expressed as the sum of prices p_ϕ on each street along the route. This allows transfers to be collected as the drivers use each street, obviating the need to know any driver's complete route. Furthermore, the price of each street only depends on the number of drivers who take that street, so the prices themselves can be determined in a decentralized fashion. We will see below that the separability of our optimal price schemes follows from the separability of each route delay function $C_i(x) = \sum_{\phi \in \Phi_i} c_\phi(u_\phi(x))$ into functions describing delays on each street along the route.²⁰

20. Under inelastic demand, the planner's information problem would not exist were it not for his inability to observe the strategy distribution x : if the planner knew x , he could easily compute the masses of active players m^r . This need not be possible if only the numbers of drivers on each street are observed. Of course, the hidden information and hidden action problems are independent when demand is elastic, since in this case even complete knowledge of the strategy distribution reveals little about the distribution of types.

4. CONGESTION PRICING UNDER INELASTIC DEMAND

Lemmas 1 and 2 guarantee that once demand is fixed, any congestion game with increasing cost functions has a unique, globally stable Nash equilibrium. Unfortunately, this equilibrium is unlikely to be efficient. We define efficiency in terms of the aggregate payoff function $\bar{F}: X \rightarrow \mathbf{R}$, given by

$$\bar{F}(x) = \sum_{i \in S} x_i F_i(x) = - \sum_{i \in S} x_i C_i(x).$$

The social planner would like to choose a price scheme that implements an efficient social choice function x^* , defined by

$$x^*(m) \in \arg \max_{x \in X_m} \bar{F}(x).$$

The price scheme creates a new game whose payoffs combine transfer payments with the costs of delay. If the scheme can be chosen in such a way that the new game admits a potential function which is proportional to aggregate payoffs, we can use Lemma 1 to show that myopic adjustment with respect to the new payoffs always leads to efficient play.

Let $\hat{F}_i: X \rightarrow \mathbf{R}$ denote the payoff functions of the new game, which are the differences between the original payoffs and the prices charged by the social planner:

$$\hat{F}_i(x) \equiv F_i(x) - P_i(x).$$

We would like the new game to admit a potential function \hat{f} which is proportional to aggregate payoffs:

$$\hat{f}(x) \equiv \kappa \bar{F}(x) \quad \text{for some } \kappa > 0.$$

Moreover, we would like to create this new game using a separable price scheme.

To determine whether such a price scheme exists, we first differentiate the second identity with respect to x_i to obtain

$$\hat{F}_i(x) = \kappa \frac{\partial}{\partial x_i} \bar{F}(x).$$

When $\kappa = 1$, this equation tells us to set an individual's total payoff from choosing an action equal to the marginal social payoff of that action. Hence, the price scheme we derive will be a form of marginal cost pricing. More generally, individuals' payoffs are set proportional to marginal social payoffs. Players may bear costs which are larger or smaller than their marginal social impact, but each always agrees with the social planner about the relative costs of any action pair.

Now, recalling that payoffs in congestion games are of the form $F_i(x) = -C_i(x) = - \sum_{\phi \in \Phi_i} c_\phi(u_\phi(x))$, and observing that

$$\frac{\partial u_\phi}{\partial x_j}(x) = \begin{cases} 1 & \text{if } \phi \in \Phi_j, \\ 0 & \text{otherwise,} \end{cases}$$

we can solve for the price scheme.

$$\begin{aligned} P_i(x) &= F_i(x) - \hat{F}_i(x) \\ &= F_i(x) - \kappa \frac{\partial}{\partial x_i} \bar{F}(x) \\ &= F_i(x) - \kappa \left(F_i(x) + \sum_{j \in S} x_j \frac{\partial F_j}{\partial x_i}(x) \right) \\ &= \kappa \left(\sum_{j \in S} x_j \frac{\partial C_j}{\partial x_i}(x) \right) + (\kappa - 1)C_i(x) \end{aligned}$$

$$\begin{aligned}
 &= \kappa \left(\sum_{j \in S} x_j \left(\sum_{\phi \in \Phi_i \cap \Phi_j} c'_\phi(u_\phi(x)) \right) \right) + (\kappa - 1) \left(\sum_{\phi \in \Phi_i} c_\phi(u_\phi(x)) \right) \\
 &= \kappa \sum_{\phi \in \Phi_i} u_\phi(x) c'_\phi(u_\phi(x)) + (\kappa - 1) \sum_{\phi \in \Phi_i} c_\phi(u_\phi(x)) \\
 &= \sum_{\phi \in \Phi_i} (\kappa u_\phi(x) c'_\phi(u_\phi(x)) + (\kappa - 1) c_\phi(u_\phi(x))). \tag{T}
 \end{aligned}$$

Equation (T) shows that because the original payoff functions F_i are separable in ϕ , the price scheme we have constructed is separable as well. To create a game whose potential function is proportional to aggregate payoffs, a social planner need only set prices for the use of each individual street ϕ , and the price of each street need only depend on the number of drivers who take that street. Thus, the price scheme can be imposed in a decentralized fashion.

Letting $\bar{\eta} = \frac{1}{\kappa} - 1 > -1$, we define the *variable price scheme* $P^{\bar{\eta}}$ by

$$P_i^{\bar{\eta}}(x) = \sum_{\phi \in \Phi_i} p_\phi^{\bar{\eta}}(u_\phi(x)),$$

where the street prices $p_\phi^{\bar{\eta}}$ are given by

$$p_\phi^{\bar{\eta}} = \frac{1}{\bar{\eta} + 1} (uc'_\phi(u) - \bar{\eta}c_\phi(u)).$$

We use the word “variable” to highlight the fact that prices in these schemes vary with current utilization levels, and to contrast these schemes with the fixed price schemes defined below.

To interpret the price schemes more easily, we define

$$\eta_\phi(u) = \frac{uc'_\phi(u)}{c_\phi(u)}$$

to be the cost elasticity of street ϕ . Then when $c_\phi(u) > 0$, we can express the street prices as

$$p_\phi^{\bar{\eta}}(u) = \frac{c_\phi(u)}{\bar{\eta} + 1} (\eta_\phi(u) - \bar{\eta}).$$

The parameter $\bar{\eta}$, which we call the *elasticity threshold*, has a simple interpretation. When the cost elasticity of a street is exactly $\bar{\eta}$, the price of that street is set to zero; at cost elasticities higher than $\bar{\eta}$, positive prices are charged; at cost elasticities lower than $\bar{\eta}$, negative prices (*i.e.* subsidies) are offered.

We now characterize behaviour under the variable price schemes.

Theorem 1. *Suppose that the cost functions satisfy $uc''_\phi(u) > -2c'_\phi(u)$ for all ϕ and u . Then each variable price scheme $P^{\bar{\eta}}$ globally implements the efficient social choice function x^* .*

Proof. The condition on the cost functions c_ϕ implies that the augmented street costs $c_\phi + p_\phi^{\bar{\eta}}$ are strictly increasing:

$$\frac{d}{du}(c_\phi(u) + p_\phi^{\bar{\eta}}(u)) = \frac{d}{du} \left(\frac{1}{\bar{\eta} + 1} (uc'_\phi(u) + c_\phi(u)) \right) = \frac{1}{\bar{\eta} + 1} (uc''_\phi(u) + 2c'_\phi(u)) > 0.$$

Hence, Lemma 2 implies that the potential function \hat{f} is strictly concave.

Now, fix any demand vector m . Lemma 1(ii) implies that the unique state which maximizes \hat{f} on X_m is globally stable on X_m under all dynamics which are admissible with respect to \hat{F} . But $\hat{f}(x) = \frac{1}{\bar{\eta}+1} \bar{F}(x)$ and $\frac{1}{\bar{\eta}+1} > 0$, so this global attractor is the efficient state $x^*(m)$. \parallel

Since the potential function of the new game is proportional to aggregate payoffs, one can show that regardless of the shapes of cost functions c_ϕ , the distribution which maximizes aggregate payoffs is a locally stable Nash equilibrium. However, this distribution need only be globally stable if the potential function is concave, which is only guaranteed if the cost functions are never too concave.

We observe that an elasticity threshold of $\bar{\eta} = 0$ (which corresponds to $\kappa = 1$) yields a generalization of marginal cost pricing. In this case, the variable prices are given by $p_\phi^0(u) = uc'_\phi(u)$: players are always charged the marginal social cost of the congestion they currently create. It is not surprising that these prices render the efficient state $x^*(m)$ an equilibrium. However, Theorem 1 goes much further: it shows that if street prices are always adjusted to represent current marginal social costs, then the efficient state is the only equilibrium, is globally stable under any admissible adjustment process, and can be implemented without knowledge of the demand vector m or the efficient state $x^*(m)$.

By employing a variable price scheme, the planner ensures that the efficient allocation $x^*(m)$ is ultimately played. Even after convergence to this state occurs, the planner still faces a hidden action problem, and must continue to use tolls to maintain efficient play. Fortunately, his knowledge of the efficient state enables him to do so using an especially simple price scheme. For each $\bar{\eta} > -1$, we define the *fixed price scheme* $\Pi^{\bar{\eta},m}$ by

$$\Pi_i^{\bar{\eta},m} = \sum_{\phi \in \Phi_i} \pi_\phi^{\bar{\eta},m}, \quad \text{where } \pi_\phi^{\bar{\eta},m} = p_\phi^{\bar{\eta}}(u_\phi(x^*(m))).$$

Under the fixed price scheme $\Pi^{\bar{\eta},m}$, the planner *always* sets the prices charged by the variable price scheme $P^{\bar{\eta}}$ at the efficient state $x^*(m)$. By employing this scheme, the planner ensures that the players will return to the efficient state after any shock to their behaviour.

Theorem 2. *Suppose that the demand vector is m , and that the planner imposes the fixed price scheme $\Pi^{\bar{\eta},m}$. Then the efficient distribution $x^*(m)$ is globally stable under any admissible dynamics V .*

Proof. Let \check{F} denote payoffs under the fixed price scheme $\Pi^{\bar{\eta},m}$: that is, $\check{F}_i(x) = F_i(x) - \Pi_i^{\bar{\eta},m}$. By the definition of $\pi_\phi^{\bar{\eta}}$, $\check{F}(x^*(m)) = \hat{F}(x^*(m))$, while by construction, $x^*(m)$ is a Nash equilibrium of the variable price game \hat{F} . Therefore, $x^*(m)$ must also be a Nash equilibrium of the fixed price game \check{F} . Moreover, since the original cost functions $c_\phi(u)$ are strictly increasing, so are the cost functions $\check{c}(u) = c_\phi(u) + \pi_\phi^{\bar{\eta},m}$ of the fixed price game \check{F} . Hence, Lemma 2 implies that the potential function of \check{F} is strictly concave. It then follows from Lemma 1 that \check{F} has a unique equilibrium, and that this equilibrium is globally stable under all dynamics which are admissible with respect to \check{F} . The equilibrium must be $x^*(m)$. \parallel

We observe that Theorem 2 does not require any assumption about the convexity of costs: when prices are fixed, that delay costs c_ϕ are increasing is enough to ensure global convergence to equilibrium.

By choosing a price scheme with an arbitrary elasticity threshold and then observing the drivers' behaviour, the planner can determine the efficient state. After doing so, the planner can select other thresholds which have interesting properties in equilibrium. For example, for any demand vector m the planner can compute the threshold

$$\bar{\eta}_{\min}(m) = \min_{\phi \in \Phi} \eta_\phi(u_\phi(x^*(m))).$$

Recall that once the threshold is fixed, streets whose cost elasticities lie above the threshold are assigned positive prices, while those with elasticities below the threshold are given negative prices. Thus, setting the threshold to $\bar{\eta}_{\min}(m)$, the lowest cost elasticity obtaining at the efficient state $x^*(m)$, yields the price scheme with the lowest non-negative equilibrium prices. As long as all facilities are sensitive to congestion, $\bar{\eta}_{\min}(m)$ will exceed zero, so these prices will be strictly less than those which obtain under marginal cost pricing.

Since each of the price schemes induces the same strategy distribution $x^*(m)$, higher prices must yield higher toll revenues. By adjusting the elasticity threshold, what range of revenues can the social planner obtain?

Corollary 1. *If the demand vector m is not zero, each point in the interval $(\bar{F}(x^*(m)), \infty)$ can be achieved as the equilibrium toll revenue through an appropriate choice of elasticity threshold.*

Proof. Let $u_\phi^* = u_\phi(x^*(m))$. Observe that if $u_\phi^* > 0$, then $\lim_{\bar{\eta} \downarrow -1} p_\phi^{\bar{\eta}}(u_\phi^*) = \infty$, $\lim_{\bar{\eta} \uparrow \infty} p_\phi^{\bar{\eta}}(u_\phi^*) = -c_\phi(u_\phi^*)$, and

$$\frac{\partial}{\partial \bar{\eta}} p_\phi^{\bar{\eta}}(u_\phi^*) = \frac{-(u_\phi^* c'_\phi(u_\phi^*) + c_\phi(u_\phi^*))}{(\bar{\eta} + 1)^2} < 0.$$

Hence, equilibrium toll revenue $R_m(\bar{\eta}) = \sum_\phi u_\phi^* p_\phi^{\bar{\eta}}(u_\phi^*)$ is a smooth, strictly decreasing function which satisfies $\lim_{\bar{\eta} \uparrow \infty} R_m(\bar{\eta}) = -\sum_\phi u_\phi^* c_\phi(u_\phi^*) = \bar{F}(x^*(m))$ and $\lim_{\bar{\eta} \downarrow -1} R_m(\bar{\eta}) = \infty$. \parallel

Since $\bar{F}(x^*(m)) < 0$, Corollary 1 implies that for each demand vector m , there is a price scheme which yields a revenue of zero: by choosing an elasticity threshold of $\bar{\eta}_z(m) \equiv R_m^{-1}(0)$, the amount of revenue generated from tolls on highly congested streets can be exactly offset by subsidies paid to drivers on less congested streets. Indeed, the global stability of the efficient strategy distribution can be achieved along with any positive level of revenue, as well as a wide range of negative revenues.²¹

Of course, Corollary 1 depends crucially on the inelasticity of demand. Generating large revenues means imposing large tolls, which we would expect to cause some drivers to abandon their commute. Fortunately, our variable price schemes are still quite effective when the demand to drive depends on the cost of the trip.

5. ELASTIC DEMAND

We now extend our model to allow for elastic demand. Elastic demand introduces a new source of inefficiency: not only may players distribute themselves over the roads inefficiently, but the players who choose to drive may not be those who would drive at the social optimum. Nevertheless, our variable price schemes continue to yield efficient play.

5.1. Congestion games

To introduce elastic demand, we suppose that while all players who choose the same route experience the same delays, different players attach different values completing their commutes. Let $D^r: [0, L^r] \rightarrow [0, M^r]$ denote the demand curve for population r , so that $D^r(v)$ is the mass

21. Since the variable and fixed price schemes charge the same prices at the efficient state $x^*(m)$, Corollary 1 describes the equilibrium revenues obtained under either class of schemes.

of commuters from population r for whom the value of the commute is at least v . We assume that each D^r is onto, differentiable, and strictly decreasing: $\frac{d}{dv}D^r(v) < 0$. These assumptions imply that the inverse demand curves $\tilde{D}^r: [0, M^r] \rightarrow [0, L^r]$ are well defined and satisfy $\frac{d}{dz}\tilde{D}^r(z) < 0$. $\tilde{D}^r(z)$ is the “ z -th highest” value of commuting among players from population r , where $z \in [0, M^r]$.

If a player with valuation v commutes via route i , his payoff is $v - C_i(x)$; if he stays home, his payoff is zero. In principle, the valuations of the commuters who are active at any moment in time could be an arbitrary subsets of the sets $[0, L^r]$. However, in order to make our dynamic analysis tractable, we make the simplifying assumption that whenever a player whose valuation is v commutes, all players in his population who have higher valuations commute as well.²²

Under this assumption, behaviour in all populations can be fully described by an element x of the state space $X = \{x \in \mathbf{R}_+^n: \sum_{i \in S^r} x_i \leq M^r \text{ for all } r \in R\}$. Since the high valuation commuters are the ones who opt to drive, if we let $x^r = \sum_{i \in S^r} x_i$ denote the number of active drivers in population r , then $\tilde{D}^r(x^r)$ is the valuation of the marginal active driver. The payoffs of this marginal driver are therefore given by the (reduced form) payoff functions $\tilde{F}_i: X \rightarrow \mathbf{R}$:

$$\tilde{F}_i(x) = \tilde{D}^r(x^r) - C_i(x) \quad \text{for } i \in S^r.$$

If $\tilde{F}_i(x) > 0$, then the marginal driver would rather commute via route i than stay home; if $\tilde{F}_i(x) < 0$, he would rather stay home.

While \tilde{F}_i represents the payoffs of the marginal active driver, it also captures all players’ relative payoffs to choosing different routes.²³ Therefore, in using the reduced form payoffs to specify our evolutionary dynamics, we capture all players’ incentives for deciding between different routes, and we capture the incentives faced by the marginal active drivers in deciding whether or not to drive.

Nash equilibria of the elastic demand game must satisfy two conditions. First, if $i \in S^r$ then $\tilde{F}_i(x) = \max_{j \in S^r} \tilde{F}_j(x)$ whenever $x_i > 0$: within each population, all routes which are used perform equally well, at least as well those which are not used. Second, we must ensure that the marginal player is indifferent between commuting and staying home. When $x^r \in (0, M^r)$, the condition we need is that $\max_{j \in S^r} \tilde{F}_j(x) = 0$. When this condition holds, the valuation of the marginal player, $\tilde{D}^r(x^r)$, is equal to the cost of taking an optimal route, $\min_{j \in S^r} C_j(x)$. Players with higher valuations opt to commute, while players with lower valuations stay home.²⁴

5.2. Evolutionary dynamics

Evolutionary dynamics under elastic demand are once again defined by vector fields $V: X \rightarrow \mathbf{R}^n$. Since the number of active players can change over time, we must replace the forward invariance condition (FI 2), which kept the population masses constant, by

$$\text{For each } r, \quad \sum_{i \in S^r} V_i(x) \leq 0 \quad \text{whenever } x^r = M^r, \quad (\text{FI } 2')$$

which ensures that the number of players in each population never exceeds the upper bound of M^r .

22. This ordering property clearly must hold in any equilibrium of the game; we require it to hold out of equilibrium as well.

23. That is, the difference between a players’ payoffs to choosing route i and choosing route j equals $(v - C_i(x)) - (v - C_j(x)) = C_j(x) - C_i(x) = \tilde{F}_i(x) - \tilde{F}_j(x)$, regardless of the player’s valuation v .

24. If $x^r = 0$, so that no players in population r commute, then equilibrium requires that $\max_{j \in S^r} \tilde{F}_j(x) \leq 0$. If $x^r = M^r$, so that all players commute, then equilibrium requires that $\max_{j \in S^r} \tilde{F}_j(x) \geq 0$.

We must also modify the positive correlation condition:

$$V(x) \cdot \vec{\tilde{F}}(x) > 0 \text{ whenever } V(x) \neq \vec{0}. \tag{PC'}$$

To see that this condition has the desired interpretation, let $V_{or}(x) \equiv -\sum_{i \in S^r} V_i(x)$ and $\vec{F}_{or}(x) \equiv 0$ denote the growth rate and payoffs of the outside option in population r . Then since $\sum_{i \in S^r \cup \{or\}} V_i(x) = 0$, we see that

$$\begin{aligned} V(x) \cdot \vec{\tilde{F}}(x) &= \sum_{r \in R} \sum_{i \in S^r} V_i(x) \tilde{F}_i(x) \\ &= \sum_{r \in R} \sum_{i \in S^r \cup \{or\}} V_i(x) \tilde{F}_i(x) \\ &= \sum_{r \in R} \left(\sum_{i \in S^r \cup \{or\}} (V_i(x) - 0) \left(\tilde{F}_i(x) - \frac{1}{n^r + 1} \sum_{j \in S^r \cup \{or\}} \tilde{F}_j(x) \right) \right) \\ &= \sum_{r \in R} (n^r + 1) \text{Cov}(V^r(x), \vec{F}^r(x)). \end{aligned}$$

Hence, (PC') still requires positive correlation between growth rates and payoffs, but now the lists of strategies include outside options, and the payoffs considered are those of the marginal active player.

Finally, we state the noncomplacency condition in terms of the reduced form payoffs $\vec{\tilde{F}}$.

$$V(x) = \vec{0} \text{ implies that } x \text{ is an equilibrium of } \vec{\tilde{F}}. \tag{NC'}$$

5.3. Potential games

The definition of potential games under elastic demand is analogous to that under inelastic demand: $\vec{\tilde{F}}$ is a potential game if it admits a potential function \tilde{f} which satisfies

$$\frac{\partial \tilde{f}}{\partial x_i}(x) = \tilde{F}_i(x) \quad \text{for all } x \in X \quad \text{and} \quad i \in S.$$

Congestion games with elastic demand are therefore potential games with the potential function²⁵

$$\tilde{f}(x) = \sum_{r \in R} \int_0^{x^r} \tilde{D}^r(z) dz - \sum_{\phi \in \Phi} \int_0^{u_\phi(x)} c_\phi(z) dz.$$

Importantly, versions our lemmas from Section 2.4 can be established in this new setting.

Lemma 3. (i) *Let $\vec{\tilde{F}}$ be an elastic demand potential game, and let V be a dynamic which is admissible with respect to $\vec{\tilde{F}}$. Then every solution trajectory of V converges to a connected set of Nash equilibria of $\vec{\tilde{F}}$.*

(ii) *If the potential function \tilde{f} of $\vec{\tilde{F}}$ is strictly concave, then the maximizer of \tilde{f} on the set X is the unique equilibrium of $\vec{\tilde{F}}$ and is globally stable under V .*

Since demand is elastic, the number of active players can vary; hence, the uniqueness and global stability results in part (ii) of the lemma are with respect to all of X rather than just the subsets X_m .

Lemma 4. *Any elastic demand congestion game with distinguishable strategy distributions and with cost functions satisfying $c'_\phi > 0$ has a strictly concave potential function.*

25. Beckmann *et al.* (1956) use an analogous potential function to characterize equilibria in their elastic demand congestion model.

Proofs. The proof of Lemma 3 is a simple extension of the proof of Lemma 1. To prove Lemma 4, let $x, y \in X$; it is enough to show that the potential function is strictly concave on the line segment joining x and y . For $\lambda \in [0, 1]$, let $z(\lambda) = \lambda x + (1 - \lambda)y$, let $z_\phi(\lambda) = u_\phi(\lambda x + (1 - \lambda)y)$, and let $z^r(\lambda) = \lambda x^r + (1 - \lambda)y^r$. Then since $z'_\phi(\lambda) \neq 0$ for some $\hat{\phi} \in \Phi$, we see that

$$\frac{d^2}{(d\lambda)^2} \tilde{f}(z(\lambda)) = \sum_{r \in R} \frac{d}{dz} \tilde{D}^r(z^r(\lambda)) ((z^r)'(\lambda))^2 - \sum_{\phi \in \Phi} c'_\phi(z_\phi(\lambda)) (z'_\phi(\lambda))^2 < 0. \quad \parallel$$

Examining the proof of Lemma 4, we see that the first component of the potential function is concave as long as the demand curves D^r are downward sloping; no other conditions on demand are required. Hence, to establish the concavity of the potential function, it is enough to place restrictions on the cost functions c_ϕ .

6. EVOLUTIONARY IMPLEMENTATION

When demand is inelastic, efficiency only requires that players distribute themselves over the streets in an optimal fashion. Under elastic demand, efficiency also requires that the right players choose to drive. The planner would like to ensure efficient behaviour by imposing a price scheme, but must do so without any information about the values players attach to completing their commutes.

Formally, we let $D = (D^1, \dots, D^{\bar{r}})$ denote a profile of demand functions, and let \mathbf{D} denote the set of all such profiles. A *social choice function* $\sigma: \mathbf{D} \rightarrow X$ maps every demand profile to a strategy distribution in X . A *price scheme* is once again a map $P: X \rightarrow \mathbf{R}^n$ which for each strategy distribution x specifies the price $P_i(x)$ which users of strategy i must pay.

The price scheme creates a new elastic demand game with the same strategies as the original game. The reduced form payoffs of the new game are denoted \widehat{F} :

$$\widehat{F}_i(D, x) = \tilde{F}_i(D, x) - P_i(x).$$

Our notation makes explicit that both reduced form payoff functions depend on the demand profile D , but that the price scheme does not. We say that the price scheme P *globally implements* the social choice function σ if for each demand profile D , the strategy distribution $\sigma(D)$ is globally stable under any dynamic V which is admissible with respect to the augmented game $\widehat{F}(D, \cdot)$.

7. CONGESTION PRICING

We measure welfare under elastic demand using the function

$$W^k(D, x) = \sum_{r \in R} \int_0^{x^r} \tilde{D}^r(z) dz + k \bar{F}(x) = \sum_{r \in R} \int_0^{x^r} \tilde{D}^r(z) dz - k \sum_{i \in S} x_i C_i(x),$$

where k is some strictly positive constant. The first component of W^k captures the benefits obtained by drivers who decide to commute; each integral measures the area under the demand curve from zero through the number of active players x^r . The second component aggregates the driving times of all active players. Of course, tolls are excluded from the calculation of welfare.

The constant k measures the ratio of costs to benefits used in measuring welfare. If $k = 1$, costs and benefits are weighted equally, and W^k equals consumer surplus. However, since driving causes pollution, a “bad” not accounted for in the payoff functions, a social planner may wish to give extra weight to costs in determining welfare. By varying k (in particular, by choosing $k > 1$),

a social planner can attach added importance to reducing total driving time, which serves as a proxy for any indirect externalities which driving creates.

The social planner would like to implement the social choice function $x^k: \mathbf{D} \rightarrow X$ defined by

$$x^k(D) \in \arg \max_{x \in X} W^k(D, x).$$

To do so, he chooses a mechanism which creates a new game \widehat{F} :

$$\widehat{F}_i(D, x) \equiv \widetilde{F}_i(D, x) - P_i(x).$$

He would like the new game to have a potential function equal to the welfare function W^k .

$$\widehat{f}(D, x) \equiv W^k(D, x).$$

It is not immediately clear that a mechanism with both of these properties exists. However, differentiating the second identity with respect to x_i yields

$$\widehat{F}_i(D, x) = \widetilde{D}^r(x^r) + k \frac{\partial}{\partial x_i} \bar{F}(x).$$

Solving for P_i , we find that

$$\begin{aligned} P_i(x) &= \widetilde{F}_i(D, x) - \widehat{F}_i(D, x) \\ &= (\widetilde{D}^r(x^r) + F_i(x)) - \left(\widetilde{D}^r(x^r) + k \frac{\partial}{\partial x_i} \bar{F}(x) \right) \\ &= F_i(x) - k \frac{\partial}{\partial x_i} \bar{F}(x). \end{aligned}$$

Since this last expression does not depend on D , prices can be chosen independently of the demand profile. In fact, the last expression is equivalent to equation (T), which defined the variable price scheme $P^{\bar{\eta}}$ (with $\bar{\eta} = \frac{1}{k} - 1$). Thus, the same separable price schemes which guarantee efficient behaviour under inelastic demand also do so under elastic demand.

Theorem 3. *Suppose that the cost functions satisfy $uc''_{\phi}(u) > -2c'_{\phi}(u)$ for all ϕ and u . Then the variable price scheme $P^{\bar{\eta}}$ with elasticity threshold $\bar{\eta} = \frac{1}{k} - 1$ globally implements the social choice function x^k .*

The proof of Theorem 3 is analogous to that of Theorem 1, with Lemmas 3 and 4 used in place of Lemmas 1 and 2. We observe that the only property of demand curves required to establish this result is that they are downward sloping.

Our analysis of the inelastic demand case showed that if the number of active players in each population is fixed, the variable price schemes lead the players to distribute themselves over the streets efficiently. This suggests that if these schemes are used under elastic demand, the players who decide to stay in the game will behave efficiently. It therefore remains to show that the optimal number of players from each population will choose to stay in the game.

Were there only one population of players, the intuition for the latter claim would be simple. Lowering the elasticity threshold $\bar{\eta}$ increases the street prices. Hence, by varying $\bar{\eta}$, the planner can choose the level of demand; since the allocation of drivers will be efficient, any weighted welfare optimum can be achieved. Unfortunately, this simple argument fails when there are multiple populations. In this case, each choice of threshold $\bar{\eta}$ yields some equilibrium demand vector $(x^1(\bar{\eta}), \dots, x^r(\bar{\eta}), \dots, x^{\bar{r}}(\bar{\eta}))$. By varying $\bar{\eta}$, we can obtain demand vectors lying on a one-dimensional path through $\mathbf{R}^{\bar{r}}_+$. It is not obvious that this path should trace out the demand vectors in $\mathbf{R}^{\bar{r}}_+$ which lead to welfare optimality.

To see why it does, let $\bar{C}(x) = -\bar{F}(x)$ equal aggregate costs, and let $\hat{C}_i(x) = C_i(x) + P_i(x) = -\hat{F}_i(x)$ denote the augmented strategy costs. We designed the variable price scheme in the inelastic demand setting to set individuals' costs proportional to marginal social costs ($\hat{C}_i(x) = k \frac{\partial}{\partial x_i} \bar{C}(x)$).

Suppose that demand is inelastic. In equilibrium, the augmented costs of all strategies i used by a single population r are equal ($\hat{C}_i(x) = c^r$ whenever $x_i > 0$). Consequently, the marginal social costs of these strategies are equal ($k \frac{\partial}{\partial x_i} \bar{C}(x) = c^r$), and so aggregate behaviour is efficient.

Now suppose that demand is elastic. In equilibrium, the augmented cost of each strategy which is used must equal the valuation of the marginal driver ($\hat{C}_i(x) = \bar{D}^r(x^r)$ whenever $x_i > 0$). It then follows from the definition of the price scheme that marginal social costs are proportional to this valuation ($k \frac{\partial}{\partial x_i} \bar{C}(x) = \bar{D}^r(x^r)$). Since the valuation of the marginal driver represents the marginal social benefit of further entry, this last equation implies welfare optimality. Finally, since the variable price schemes yield games with concave potential functions, each scheme guarantees the global stability of the welfare maximizing state.

We can also explain our price schemes directly in terms of the potential functions. In the inelastic demand setting, we chose the variable prices so that the potential function \hat{f} was equal to k times aggregate payoffs, so that admissible dynamics would ascend the aggregate payoff function. Under elastic demand, the potential function \hat{f} and the welfare function W^k both have a new term capturing players' valuations. But in each case, this term is $\sum_r \int_0^{x^r} \bar{D}^r$. Therefore, the variable prices which equate \hat{f} and $k\bar{F}$ also equate \hat{f} and W^k , and so ensure welfare optimality.

More intuitively, a social planner is able to set the relative costs of playing the various strategies appropriately while fixing the absolute costs at any desired scale. If the social planner chooses this scale correctly, then a player comparing his benefit from joining the game to his cost of participating will decide to join precisely when it is socially desirable for him to do so.

After imposing the variable price scheme, the planner can learn the efficient state $x^*(D)$ by observing the players' behaviour. Since behaviour is anonymous, maintaining efficient play also requires the use of a price scheme. As before, this maintenance can be accomplished by administering a fixed price scheme. For each demand vector D , define the fixed price scheme $\Pi^{\bar{\eta}, D}$ by

$$\Pi_i^{\bar{\eta}, D} = \sum_{\phi \in \Phi_i} \pi_{\phi}^{\bar{\eta}, D}, \quad \text{where } \pi_{\phi}^{\bar{\eta}, D} = p_{\phi}^{\bar{\eta}}(u_{\phi}(x^k(D))) \quad \text{and} \quad k = \frac{1}{\bar{\eta} + 1}.$$

A variation on the proof of Theorem 2 establishes the following result.

Theorem 4. *Suppose that the demand profile is D , and that the planner imposes the fixed price scheme $\Pi^{\bar{\eta}, D}$ with $\bar{\eta} = \frac{1}{k} - 1$. Then the efficient distribution $x^k(D)$ is globally stable under any admissible dynamics V .*

Again, no convexity condition on street costs c_{ϕ} is required.

Under elastic demand, the central role of the choice of $\bar{\eta} \in (-1, \infty)$ is to allow the implementation of efficient behaviour under different welfare measures; hence, different choices of $\bar{\eta}$ generate different equilibrium behaviours. Nevertheless, we can give some indication of the equilibrium toll revenues generated by different choices of $\bar{\eta}$.²⁶ For $\bar{\eta}$ near -1 (and hence $k = \frac{1}{\bar{\eta} + 1}$ very large), toll revenue is zero: prices are so high that no one commutes. Whenever

26. For simplicity, we assume here that delay costs satisfy $c_{\phi}(0) > 0$ for all streets $\phi \in \Phi$.

$\bar{\eta} \leq 0$, all prices are weakly positive, so toll revenues are weakly positive as well. Similarly, when $\bar{\eta} \geq \eta_{\max} = \max_{\phi, u} \eta_{\phi}(u)$, prices and revenues are weakly negative. A continuity argument then shows that for any demand profile D , there is an elasticity threshold $\bar{\eta}_z(D) \rightarrow [0, \eta_{\max}]$ which is revenue neutral in equilibrium. This threshold typically corresponds to a welfare measure with $k < 1$, which places relatively little weight on keeping aggregate delays small.

8. DISCUSSION

We considered a problem faced by a planner who would like to ensure the efficient use of a highway network. The problem entails both hidden information and hidden actions. We showed that the planner can guarantee efficient behaviour by introducing a separable, variable price scheme. After using this scheme to resolve his information problem, the planner can maintain efficient behaviour using a separable, fixed price scheme.

Throughout our analysis, we have assumed that the planner is fully patient, only caring about the final outcomes of the evolutionary process. Alternatively, one could consider a less patient planner who cares about the entire time path of play. For example, the planner might evaluate the time path of play by aggregating the welfare of behaviour over time using some fixed discount rate. To study this issue, one would need to introduce conditions on dynamics specifying the speed at which behaviour adjustment takes place; for instance, one could assume that the difference between strategies' growth rates is monotone in the difference between their payoffs. In this case, price schemes that dramatically penalize the use of crowded routes would lead to faster adjustment towards the efficient state. If demand is inelastic, this would favour the use of schemes with low elasticity thresholds $\bar{\eta}$, as such schemes magnify payoff differences between routes. However, low values of $\bar{\eta}$ also generate high tolls; if demand is elastic, such tolls might reduce network usage to undesirably low levels. The planner might strike a compromise between these issues by judiciously adjusting $\bar{\eta}$ over time.

Our study of congestion pricing has focused on traffic during a single usage period. An important use of congestion pricing which we do not address here is that of smoothing peak demand over different usage periods: that is, setting prices to encourage driving before and after the period of peak demand.²⁷ Allowing commuters the choice of when to drive introduces new modelling issues which we hope to address in future research.

While this paper has focused on pricing roadway networks, the techniques developed here seem applicable in principle to the pricing of computer networks.²⁸ Like roadway networks, computer networks are used by large numbers of individuals who are primarily concerned with the speed of their own conveyance through the network. On the other hand, there are some important differences between the two types of networks. For one, routing decisions in computer networks are presently made by routing devices positioned throughout the network rather than by individual users. In addition, different network applications require vastly different levels of service: while a real-time video conference requires a constant, high bandwidth connection throughout its duration, e-mail messages require little bandwidth, and can be sent after some delay without significantly reducing user benefits. For these reasons, a realistic model of computer network congestion would be significantly different from the model considered here. Nevertheless, evolutionary implementation in computer networks seems an important topic for future research.

27. This aspect of congestion pricing is considered by Vickrey (1963, 1969) and Arnott *et al.* (1990, 1993).

28. Congestion pricing in computer networks is studied by Cocchi *et al.* (1993), Mackie-Mason and Varian (1995), Shenker *et al.* (1996) and Crémer and Hariton (1999).

APPENDIX

The Proof of Lemma 1

Suppose that the dynamic V is admissible with respect to the game F . Since the potential function f is a global Lyapunov function for this dynamic (i.e. since $\frac{d}{dt} f(x_t) = \nabla f(x_t) \cdot \dot{x}_t = F(x_t) \cdot V(x_t) \geq 0$), Proposition 1 of Losert and Akin (1983) and Theorem 5.4.1 of Robinson (1995) imply that every solution trajectory of the dynamic converges to a connected set of rest points. Condition (NC) tells us that all rest points of the dynamic are Nash equilibria of F , proving part (i) of the lemma.

The proof of part (ii) requires the following preliminary results.

Lemma A1. *Suppose that the dynamic V is admissible with respect to the game F . Then x^* is a rest point of the dynamic if and only if it is a Nash equilibrium of F .*

Proof. The “only if” direction follows directly from condition (NC), so we need only consider the “if” direction. Let x be a Nash equilibrium of F . Let $S_d^r(x)$ be the set of strategies in S^r that are in decline at x : $S_d^r(x) = \{i \in S^r : V_i(x) < 0\}$. Since $S_d^r(x)$ can only contain strategies which are used at x (by condition (FI 1)), and since all such strategies must be optimal, we see that $S_d^r(x) \subset \arg \max_{j \in S^r} F_j(x)$. But $\sum_{i \in S^r} V_i(x) = 0$ by the forward invariance condition (FI 2). Thus, the inclusion implies that $\sum_{r \in R} \sum_{i \in S^r} V_i(x) F_i(x) \leq 0$. Summing over r , we see that $\sum_{r \in R} \sum_{i \in S^r} V_i(x) F_i(x) \leq 0$. We therefore conclude from condition (PC) that $V(x) = \vec{0}$. \parallel

Lemma A2. *Suppose that F is a potential game with potential function f . Then x is a Nash equilibrium of F if and only if it satisfies the Kuhn–Tucker first-order conditions for a maximizer of f on X_m : that is, for some $\mu \in \mathbf{R}^R$ and some $\lambda \in \mathbf{R}^S$,*

$$\frac{\partial f x}{\partial x_i} = \mu^r - \lambda_i \text{ for all } i \in S^r \text{ and } r \in R, \tag{KT1}$$

$$\lambda_i x_i = 0 \text{ for all } i \in S, \text{ and} \tag{KT2}$$

$$\lambda_i \geq 0 \text{ for all } i \in S. \tag{KT3}$$

Proof. Suppose that x is a Nash equilibrium of F . Since $F_i(x) = \frac{\partial f}{\partial x_i}(x)$, (KT1)–(KT3) are satisfied by x , $\mu^r = \max_{j \in S^r} F_j(x)$, and $\lambda_i = \mu^r - F_i(x)$ for $i \in S^r$. The proof of the converse is similar. \parallel

Now suppose that f is strictly concave. Then f has a unique maximizer x^* on X_m , which is the unique point satisfying (KT1)–(KT3) (see, e.g. Theorems 4.38 and 4.39 of Avriel, 1976). Thus, Lemma A2 implies that x^* is the unique Nash equilibrium of F , and so Lemma A1 implies that x^* is the unique rest point of $\dot{x} = V(x)$. Hence, part (i) of the lemma implies that x^* is a global attractor under $\dot{x} = V(x)$. \parallel

We now consider evolution under *payoff monotone* dynamics. Payoff monotone dynamics are dynamics of the form

$$\dot{x}_i = x_i g_i(x) \quad \text{for all } i \in S^r \text{ and } r \in R, \tag{PM}$$

where the percentage growth rates $g_i : X \rightarrow \mathbf{R}$ satisfy the monotonicity condition

$$g_i(x) > g_j(x) \quad \text{if and only if } F_i(x) > F_j(x)$$

for all populations $r \in R$ and strategies $i, j \in S^r$. In addition, the functions g_i are required to be Lipschitz continuous (which implies condition (LC)), and must satisfy $\sum_{i \in S^r} x_i g_i(x) = 0$ for all populations $r \in R$ (which is condition (FI 2)).

An important example of a payoff monotone dynamic is the *replicator dynamic*, which is defined by

$$\dot{x}_i = x_i \left(F_i(x) - \frac{1}{m^r} \sum_{j \in S^r} x_j F_j(x) \right) \quad \text{for all } i \in S^r \quad \text{and} \quad r \in R. \tag{R}$$

In words, the replicator dynamic requires that the percentage growth rate of each strategy is given by the difference between that strategy’s payoff and the average payoff in its population. This dynamic can be interpreted as a model of evolution through imitation—see Björnerstedt and Weibull (1996) and Schlag (1998).

It is easy to see that in addition to satisfying conditions (LC) and (FI 2), payoff monotone dynamics also satisfy condition (FI 1). We now show that in these dynamics also satisfy the positive correlation condition (PC).

Lemma A3. *Every payoff monotone dynamic satisfies condition (PC).*

Proof. Let (PM) be a payoff monotone dynamic, and suppose that $x \in X$ is not a rest point of (PM). Then the set of populations which are not at rest at x , $\tilde{R} = \{r \in R : x_i g_i(x) \neq 0 \text{ for some } i \in S^r\}$, is nonempty. Condition (FI 2) then implies that for all populations $r \in \tilde{R}$, the sets of strategies $G^r = \{i \in S^r : g_i(x) > 0\}$ and $S^r - G^r$ are nonempty. Thus, for such populations, the scalars $l^r = \min_{i \in G^r} F_i(x)$ and $h^r = \max_{i \in S^r - G^r} F_i(x)$ are well defined, and, by payoff monotonicity, satisfy $l^r > h^r$. Therefore, since $\sum_{i \in G^r} x_i g_i(x) = -\sum_{i \in S^r - G^r} x_i g_i(x)$ by condition (FI 2), we find that

$$\begin{aligned} V(x) \cdot F(x) &= \sum_{r \in R} \sum_{i \in S^r} (x_i g_i(x)) F_i(x) \\ &\geq \sum_{r \in \tilde{R}} \left(l^r \sum_{i \in G^r} x_i g_i(x) + h^r \sum_{i \in S^r - G^r} x_i g_i(x) \right) \\ &= \sum_{r \in \tilde{R}} \left((l^r - h^r) \sum_{i \in G^r} x_i g_i(x) \right) > 0. \quad \parallel \end{aligned}$$

Payoff monotone dynamics admit many rest points which are not Nash equilibria, and therefore fail condition (NC). To describe the rest points of these dynamics, we define the set of *restricted equilibria* of the game F by

$$RE = \{x \in X_m : [x_i > 0 \Rightarrow F_i(x) = \max_{j \in S^r : x_j > 0} F_j(x)] \quad \text{for all } i \in S^r \quad \text{and} \quad r \in R\}.$$

Thus, x is a restricted equilibrium of F if it is a Nash equilibrium of a restricted version of F in which only strategies in the support of x can be played. (For example, every pure strategy profile is a restricted equilibrium.) It is easily verified that the rest points of a payoff monotone dynamic are precisely the restricted equilibria of the underlying game. Nevertheless, if the game F has a strictly concave potential function, we can establish the following result.

Proposition A1. *Suppose that the game F admits a strictly concave potential function f . Then under any payoff monotone dynamic, any solution trajectory starting from an interior initial condition converges to the unique Nash equilibrium of F .*

Proof. We begin by showing that the set of restricted equilibria of F , and hence the set of rest points of the payoff monotone dynamic, is finite. To see this, consider any restricted version of F in which each population is constrained to choose from some nonempty subset of its pure strategies from F . Then by definition, the original potential function f also serves as potential function for this restricted game. Since f is strictly concave, Lemma A2 and the argument which follows it show that the restricted game has a unique Nash equilibrium, which is the point which maximizes f on the appropriate subset of X_m . This Nash equilibrium is a restricted equilibrium of the original game F , and all restricted equilibria are of this form. Since the number of restricted games which can be derived from F is finite, so too is the number of restricted equilibria.

Now, since the dynamic satisfies condition (PC) by Lemma A3, it admits the global Lyapunov function f . Thus, Proposition 1 of Losert and Akin (1983) and Theorem 5.4.1 of Robinson (1995) imply that every solution trajectory converges to a connected set of rest points. Since the number of rest points is finite, convergence must always be to a unique limit point, which is a restricted equilibrium.

Let $\{x_t\}_{t \geq 0}$ be a solution trajectory of the dynamic with interior initial condition x_0 ; then x_t lies in the interior of X_m for all finite times t (see, e.g. Weibull, 1995, p. 195). Let $y = \lim_{t \rightarrow \infty} x_t$, and suppose that y is a restricted equilibrium of F which is not a Nash equilibrium of F . Then for some population $r \in R$, there is a strategy $i \in S^r$ such that $y_i = 0$ and such that $F_i(y) > F_j(y)$ for each $j \in S^r$ in the support of y . It then follows from payoff monotonicity and the fact that y is a rest point that $g_i(y) > g_j(y) = 0$. Since the function g_i is continuous, it follows that for some $\varepsilon > 0$, there exists a $\delta > 0$ such that $g_i(x) > \varepsilon$ whenever $|x - y| < \delta$. Now since $\{x_t\}_{t \geq 0}$ converges to y , there exists a T such that $|x_t - y| < \delta$ whenever $t \geq T$. At such times t , $(\dot{x}_t)_i > (x_t)_i \varepsilon > 0$. Hence, $(x_t)_i = (x_T)_i + \int_T^t (\dot{x}_s)_i ds \geq (x_T)_i > 0$ for all $t \geq T$, contradicting the definition of the trajectory. We therefore conclude that $\{x_t\}_{t \geq 0}$ approaches the unique Nash equilibrium of F . \parallel

Acknowledgements. I thank Eddie Dekel, James Dow, Jeff Ely, Bart Lipman, Nolan Miller, Larry Samuelson, Jeroen Swinkels, three anonymous referees, and seminar audiences at Wisconsin, the Royal Institute of Technology (Stockholm), the Game Theory Society World Congress (Bilbao) and the Econometric Society World Congress (Seattle) for their comments. Early versions of some results in this paper can be found in Chapter 4 of my doctoral dissertation (Sandholm, 1998). Financial support from a State Farm Companies Foundation Dissertation Fellowship and from NSF Grant SES-0092145 is gratefully acknowledged.

REFERENCES

ABREU, D. and MATSUSHIMA, H. (1994), "Exact Implementation", *Journal of Economic Theory*, **64**, 1–19.
 ARNOTT, R., DE PALMA, A. and LINDSEY, R. (1990), "Economics of a Bottleneck", *Journal of Urban Economics*, **27**, 111–130.

- ARNOTT, R., DE PALMA, A. and LINDSEY, R. (1993), "A Structural Model of Peak Period Congestion: A Traffic Bottleneck with Elastic Demand", *American Economic Review*, **83**, 161–179.
- AUMANN, R. and BRANDENBURGER, A. (1995), "Epistemic Conditions for Nash Equilibrium", *Econometrica*, **63**, 1161–1180.
- AVRIEL, M. (1976), *Nonlinear Programming: Analysis and Methods* (Englewood Cliffs, NJ: Prentice-Hall).
- BECKMANN, M., MCGUIRE, C. B. and WINSTEN, C. B. (1956), *Studies in the Economics of Transportation* (New Haven: Yale University Press).
- BERGER, U. and HOFBAUER, J. (2000), "The Nash Dynamics" (Mimeo, Universität Wien).
- BJÖRNERSTEDT, J. and WEIBULL, J. W. (1996), "Nash Equilibrium and Evolution by Imitation", in K. J. Arrow *et al.* (eds.) *The Rational Foundations of Economic Behavior* (New York: St. Martin's Press).
- BROWN, G. W. and VON NEUMANN, J. (1950), "Solutions of Games by Differential Equations", in H. W. Kuhn and A. W. Tucker (eds.) *Contributions to the Theory of Games I*, Annals of Mathematics Studies 24 (Princeton: Princeton University Press).
- CABRALES, A. (1999), "Adaptive Dynamics and the Implementation Problem with Complete Information", *Journal of Economic Theory*, **86**, 159–184.
- CABRALES, A. and PONTI, G. (2000), "Implementation, Elimination of Weakly Dominated Strategies and Evolutionary Dynamics", *Review of Economic Dynamics*, **3**, 247–282.
- CLARKE, E. (1971), "Multipart Pricing of Public Goods", *Public Choice*, **2**, 19–33.
- COCCHI, R., SHENKER, S., ESTRIN, D. and ZHANG, L. (1993), "Pricing in Computer Networks: Motivation, Formulation, and Example", *IEEE/ACM Transactions on Networking*, **1**, 614–627.
- CRÉMER, J. and HARITON, C. (1999), "The Pricing of Critical Applications in the Internet" (Mimeo, Université de Toulouse 1).
- DE TRENQUALYE, P. (1988), "Stability of the Groves and Ledyard Mechanism", *Journal of Economic Theory*, **46**, 164–171.
- DE TRENQUALYE, P. (1989), "Stable Implementation of Lindahl Allocations", *Economics Letters*, **29**, 291–294.
- GLAZER, J. and MA, C.-T. A. (1989), "Efficient Allocation of a 'Prize'—King Solomon's Dilemma", *Games and Economic Behavior*, **1**, 222–233.
- GOMEZ-IBÁÑEZ, J. A. and SMALL, K. A. (1994), *Road Pricing for Congestion Management: A Survey of International Practice* (Washington: National Academy Press).
- GROVES, T. (1973), "Incentives in Teams", *Econometrica*, **41**, 617–631.
- JORDAN, J. S. (1986), "Instability in the Implementation of Walrasian Allocations", *Journal of Economic Theory*, **39**, 301–328.
- LOSERT, V. and AKIN, E. (1983), "Dynamics of Games and Genes: Discrete versus Continuous Time", *Journal of Mathematical Biology*, **17**, 241–251.
- MACKIE-MASON, J. K. and VARIAN, H. (1995), "Pricing the Internet", in K. Brian and K. James (eds.) *Public Access to the Internet* (Cambridge, MA: MIT Press).
- MONDERER, D. and SHAPLEY, L. S. (1996), "Potential Games", *Games and Economic Behavior*, **14**, 124–143.
- MUENCH, T. and WALKER, M. (1983), "Are Groves–Ledyard Equilibria Attainable?", *Review of Economic Studies*, **50**, 393–396.
- MYERSON, R. B. (1982), "Optimal Coordination in Generalized Principal-Agent Problems", *Journal of Mathematical Economics*, **10**, 67–81.
- PHANG, S.-Y. and TOH, R. S. (1997), "From Manual to Electronic Road Congestion Pricing: The Singapore Experience and Experiment", *Transportation Research E*, **33**, 97–106.
- PONTI, G. (2000), "Splitting the Baby in Two: Solving Solomon's Dilemma with Boundedly Rational Agents", *Journal of Evolutionary Economics*, **10**, 449–455.
- REPULLO, R. (1987), "A Simple Proof of Maskin's Theorem on Nash Implementation", *Social Choice and Welfare*, **4**, 39–41.
- ROBINSON, C. (1995), *Dynamical Systems* (Boca Raton, FL: CRC Press).
- ROSENTHAL, R. W. (1973), "A Class of Games Possessing Pure Strategy Nash Equilibria", *International Journal of Game Theory*, **2**, 65–67.
- SANDHOLM, W. H. (1998), "Evolutionary Justification of Nash Equilibrium" (Ph.D. Dissertation, Northwestern University).
- SANDHOLM, W. H. (2001), "Potential Games with Continuous Player Sets", *Journal of Economic Theory*, **97**, 81–108.
- SCHLAG, K. H. (1998), "Why Imitate, and If So, How? A Boundedly Rational Approach to Multi-armed Bandits", *Journal of Economic Theory*, **78**, 130–156.
- SCHRANK, D. and LOMAX, T. (1999), "The 1999 Annual Mobility Report" (Mimeo, Texas Transportation Institute, Texas A & M University).
- SHENKER, S., CLARK, D., ESTRIN, D. and HERZOG, S. (1996), "Pricing in Computer Networks: Reshaping the Research Agenda", *Telecommunications Policy*, **20**, 183–201.
- SJÖSTRÖM, T. (1994), "Implementation in Undominated Nash Equilibrium without Integer Games", *Games and Economic Behavior*, **6**, 502–511.
- SMALL, K. and GOMEZ-IBÁÑEZ, J. A. (1998), "Road Pricing for Congestion Management: The Transition from Theory to Policy", in K. J. Button and E. T. Verhoef (eds.) *Road Pricing, Traffic Congestion, and the Environment* (Cheltenham, UK: Edward Elgar).
- SOO, C. (1998), "ERP Cuts Congestion in Singapore", *Traffic Technology International*, June/July, 23–24.
- The Economist* (1997), "Living with the Car", December 6–12, 21–23.

- VEGA-REDONDO, F. (1989), "Implementation of Lindahl Equilibrium: An Integration of the Static and Dynamic Approaches", *Mathematical Social Sciences*, **18**, 211–228.
- VICKREY, W. (1961), "Counterspeculation, Auctions, and Competitive Sealed Tenders", *Journal of Finance*, **16**, 8–37.
- VICKREY, W. (1963), "Pricing in Urban and Suburban Transport", *American Economic Review Papers and Proceedings*, **52**, 452–465.
- VICKREY, W. (1969), "Congestion Theory and Transport Investment", *American Economic Review Papers and Proceedings*, **59**, 251–260.
- WALKER, M. (1984), "A Simple Auctioneerless Mechanism with Walrasian Properties", *Journal of Economic Theory*, **32**, 111–127.
- WEIBULL, J. W. (1995), *Evolutionary Game Theory* (Cambridge, MA: MIT Press).
- WEIBULL, J. W. (1996), "The Work of John Nash in Game Theory", *Journal of Economic Theory*, **69**, 153–185.