Annals of Operations Research 113, 15–26, 2002

Nash Equilibrium and Subgame Perfection in Observable Queues

REFAEL HASSIN hassin@post.tau.ac.il
Department of Statistics and Operations Research, Tel Aviv University, Tel Aviv 69978, Israel

MOSHE HAVIV haviv@mscc.huji.ac.il
Department of Statistics, The Hebrew University 91905 Jerusalem, Israel and Department of Econometrics and Business Statistics, The University of Sydney, Sydney, NSW 2006, Australia

Abstract. A subgame perfection refinement of Nash equilibrium is suggested for games of the following type: each of an infinite number of identical players selects an action using his private information on the system's state; any symmetric strategy results in a discrete Markov chain over such states; the player's payoff is a function of the state, the selected action, and the common strategy selected by the other players. The distinction between equilibria which are subgame perfect and those which are not, is made apparent due to the possibility that some states are transient. We illustrate the concept by considering several queueing models in which the number of customers in the system constitutes the state of the system.

1. Introduction

The solution concept of Nash equilibrium is commonly used in models of queueing systems (see [8] for many examples), and in many cases multiple equilibria exist. This is a common phenomenon in *observable queues*, that is, systems in which an arriving customer observes the queue before making a decision. A possible refinement of the equilibrium concept is that of subgame perfection. It prescribes optimal responses for every state including those which are out of the equilibrium path, i.e., states that will not be visited while the strategy under consideration is used.

We consider models with infinitely many identical decision makers (players), each facing a state from a state-space S. Each state $s \in S$ is associated with a set of actions $A(s)$. A pure strategy δ, specifies an action, $\delta(s) \in A(s)$, for every $s \in S$. Randomized (mixed) strategies are also allowed, and such strategies define a set of lotteries on action sets, a lottery per state. The probability of selecting action $a \in A(s)$ by strategy δ when observing state $s \in S$ is denoted by $\delta(a|s)$. The set of all (pure and randomized) strategies is denoted by Δ. The use of a symmetric strategy in which all players adopt the same strategy, results in a time homogeneous Markov process over S. We denote this process by $\{X_t\}$, $t \geqslant 0$. A symmetric strategy coupled with an initial state determines the process statistically.

Example 1 (FCFS queues with balking). This example is based on the model of Naor [9]. Consider a first-come first-served (FCFS) single server memoryless queue, with an arrival rate λ and a service rate μ. Let X_t, $t \geqslant 1$, be the number of customers in the system at the time of arrival of the tth customer. Thus, $S = \{0, 1, 2, \ldots\}$. The initial state may be any state in S. A customer observes upon his arrival the queue length $s \in S$, and selects one action out of two: leave for good ("action 0", usually referred to as *balk*) or join ("action 1"). The set of joining probabilities $\delta(1|s)$, $s \in S$, completely describes the strategy.

To every triple (s, a, δ) of $s \in S$, $a \in A(s)$ and $\delta \in \Delta$, corresponds a real value denoted by $R_s(a, \delta)$. This value is the (expected) payoff for a player who observes state $s \in S$ and selects action $a \in A(s)$, when all others use strategy δ.

Example 1 (continued). Suppose that the value of service to a customer is R, and that each unit of time in the system costs him C, where $R > C/\mu$. Then, $R_s(0, \delta) = 0$ and $R_s(1, \delta) = R - (s + 1)/\mu C$.

This payoff reflects the assumption that a customer may balk at no cost, whereas if he joins, he receives R after staying in the system during a time interval of expected length $(s + 1)/\mu$. In this example, the value of $R_s(a, \delta)$ is not a function of δ.

Definition 1. For a strategy δ and a state $s \in S$, action $a \in A(s)$ is an *optimal response*, if

$$a \in \arg\max_{a' \in A(s)} R_s\big(a', \delta\big).$$

A pure strategy γ is a *dominant strategy* if for any $s \in S$ and any $\delta \in \Delta$,

$$\gamma(s) \in \arg\max_{a' \in A(s)} R_s(a, \delta).$$

Dominance requires that one strategy is the best option for a player *regardless* of the strategy selected by others. Dominant strategies usually exist only in trivial cases. Therefore, there is a need for a less restrictive solution concept.

Let $\pi_s(\delta)$ be the limit probability (when $t \to \infty$) of $X_t = s$, given that $X_0 = s$ and that strategy δ is used by all. This probability is assumed to exist, and in fact all that is required is to assume aperiodicity in the underlying Markov process. Since the existence of a single ergodic class is not guaranteed, this limit probability is a function of the initial state.[1]

For a given strategy δ, a state s is *recurrent* if $\pi_s(\delta) > 0$.[2] A recurrent state s with $\pi_s(\delta) = 1$ is *absorbing*. A non-recurrent state s, i.e., a state s with $\pi_s(\delta) = 0$, is called *transient*.[3]

[1] Using the same initial and target states while defining such limit probabilities is standard. It is possible, however, to define the limit differently, for example, when the initial state is fixed to some focal state, like $s = 0$ in example 1. This may lead to different results but we do not discuss them.

[2] This terminology is not standard, usually such states are called *positive recurrent*.

[3] Again, this is not a standard terminology. A state s with $\pi_s(\delta) = 0$ but with a probability 1 of ever returning to it, is usually called *null recurrent*.

Example 1 (continued). For a strategy δ, let $n(\delta) = \inf\{s \mid \delta(1|s) = 0\}$. When $n(\delta)$ is finite, the only recurrent states under δ are $s = 0, 1, \ldots, n(\delta)$, and they compose a single recurrent class. In other words, $\pi_s(\delta) > 0$ for $s = 0, \ldots, n(\delta)$ and $\pi_s(\delta) = 0$ for $s > n(\delta)$. When δ is the strategy of never balking, $n(\delta) = \infty$. In this case, $\pi_s(\delta) > 0$ for all $s \geqslant 0$ if and only if $\sum_{n=1}^{\infty} \rho^n \prod_{i=1}^{n-1} \delta(1|s) < \infty$, where $\rho = \lambda/\mu$. Otherwise, $\pi_s(\delta) = 0$ for all $s \geqslant 0$.

Definition 2. A strategy δ is an *equilibrium strategy* if it maximizes the player's expected payoff $\sum_{s\in S} \pi_s(\delta) R_s(a, \delta)$, given that the other players adopt it. Equivalently, $\delta(a|s) > 0$ only if $a \in \arg\max_{a' \in A(s)} R_s(a', \delta)$ for all $s \in S$ with $\pi_s(\delta) > 0$.

The definition assumes that steady-state conditions have been reached, and in particular, the state in which the process initiates belongs to the ergodic class that contains the observed state. Thus, δ is an equilibrium strategy if it prescribes optimal response actions *for all states which are recurrent under it.*

Remark 1. The condition for equilibrium does not impose restrictions on the actions prescribed for transient states. However, this does not mean that prescribing arbitrary actions to these states necessarily leads to another equilibrium. This, in part, is due to the fact that under a new action the transient state may become recurrent.

Example 1 (continued). When $\lambda \geqslant \mu$, "always join" is an equilibrium strategy, since all states are transient under this strategy.

Let

$$n_e = \max\left\{s\colon\ R - \frac{s+1}{\mu} C \geqslant 0\right\}.$$

Any *pure* strategy δ with $n(\delta) = n_e$ is an equilibrium. These strategies have the same set of recurrent states and they only differ in the actions they prescribe to transient states. This includes strategies such as $\delta(1|s) = 1$ for all $s \neq n_e$ and $\delta(1|n_e) = 0$. Note that this strategy prescribes joining long queues. Yet, states s with $s > n_e$ are transient under it. Assuming $R - (n_e + 1)/\mu C > 0$, the conditions $\delta(1|s) = 1$ for $0 \leqslant s \leqslant n_e - 1$ and $\delta(1|n_e) = 0$ are necessary and sufficient for any strategy to be an equilibrium.[4]

Example 1 is typical in the sense that it possesses multiple equilibria. A possible refinement of the equilibrium concept is given next.

[4] If $R - (n_e + 1)/\mu C = 0$, the conditions are that for some p, $0 \leqslant p \leqslant 1$

$$\delta(1|s) = \begin{cases} 1, & 0 \leqslant s \leqslant n_e - 2, \\ p, & s = n_e - 1, \\ 0, & s = n_e. \end{cases}$$

If $p = 0$, the requirement $\delta(1|n_e) = 0$ can be removed.

Definition 3. A strategy δ is a *subgame-perfect equilibrium strategy* (SPE) if $a \in \arg\max_{a' \in A(s)} R_s(a', \delta)$ for every $s \in S$ and $a \in A(s)$ such that $\delta(a|s) > 0$.

A SPE strategy is an equilibrium strategy, but the converse is not necessarily true. The additional requirement from a strategy to be a SPE, once it is established that it is an equilibrium, is that it prescribes optimal response actions also at states that will not be encountered by a player when all follow the equilibrium strategy.

Example 1 (continued). Assume that $R - (n_e + 1)/\mu C > 0$, then the unique SPE strategy is $\delta(1|s) = 1$ for $0 \leqslant s \leqslant n_e - 1$ and $\delta(1|s) = 0$ for $s \geqslant n_e$.[5]

2. Examples

In this section we discuss four other examples. Examples 2, and 3 are similar to example 1 except for the service disciplines. In example 4, customers decide whether to purchase the right to belong to a priority class of customers. In example 5 customers decide whether to use a shuttle that departs as soon as it is full or a bus that departs at random times. In all these situations, arrivals are aware of the state of the system, i.e., we deal with observable queues.

2.1. *Example 2: Processor sharing*

This model is similar to the one of example 1 except for that the service discipline is *egalitarian processor sharing* (EPS). This means that the server's capacity is evenly shared by those present in the system. In particular, when $n \geqslant 1$ customers are present, each of them is equally likely, with probability $(\mu/n)\Delta t + o(\Delta t)$, to complete service during the next Δt unit of times. Reneging after joining is not allowed.[6]

The states, actions, strategies, transition probabilities and $R_s(0, \delta) = 0$ are as in example 1. The difference is in $R_s(1, \delta)$, which is defined here as

$$R_s(1, \delta) = R - CE(s, \delta), \quad s \geqslant 0, \tag{1}$$

[5] For $R - (n_e + 1)/\mu C = 0$, the uniqueness property is lost and the SPE strategies have the form

$$\delta(1|s) = \begin{cases} 1, & 0 \leqslant s \leqslant n_e - 2, \\ p, & s = n_e - 1, \\ 0, & s \geqslant n_e. \end{cases}$$

[6] In example 1, the lack of reneging in equilibrium can be derived from the model as one's fate is independent of later arrivals. This is not the case in the EPS. See [3] for details. In the current example, reneging is prohibited as an assumption of the model.

where $E(s, \delta)$ is the expected time in the system for a customer when there $s \geqslant 0$ other customers in the system and strategy δ is used by all.[7]

This model was considered by Altman and Shimkin [2]. They showed that there exists a unique SPE strategy (they refer to it as an equilibrium). This strategy is based on a threshold. A strategy with threshold $x = n + r$ where $n = \lfloor x \rfloor$ and $r \geqslant 0$, prescribes joining whenever the number of customers in the system is in $\{0, \ldots, n-1\}$, balking whenever it is in $\{n+1, n+2, \ldots\}$, and joining with probability r when it is n. An algorithm which determines the unique SPE strategy is given in [2].

If the unique SPE strategy is pure ($r = 0$), then any strategy which prescribes joining at $\{0, \ldots, n-1\}$ and balking at n (regardless of what it prescribes elsewhere), is an equilibrium. When the unique SPE strategy is mixed with threshold $x = n + r$, then any strategy which prescribes joining at $\{0, \ldots, n-1\}$, joining with probability r at n, and balking at $n + 1$ is an equilibrium.

2.2. *Example 3: LCFS-PR without reneging*

This example is similar to example 1 except for that the service discipline is *last-come first-served with preemption-resume* (LCFS-PR). Specifically, a customer commences service as soon as he arrives, possibly preempting the service of another customer. Customers return to service in a reversed order of their arrival times. Resumed service takes place from the point of its interruption. Reneging is not allowed.[8] The model was introduced by Tilt and Balachandran [10] who also observed that the equilibrium is not unique. However, the solutions obtained there are not SPE.[9]

The states, actions, strategies, transition probabilities and $R_s(0, \delta) = 0$ are as in example 1. The difference is in $R_s(1, \delta)$ which is defined here to be

$$R_s(1, \delta) = R - CE(s, \delta), \quad s \geqslant 0, \tag{2}$$

where $E(s, \delta)$ is the expected time in the system for one *who is in service* when an additional number of $s \geqslant 0$ customers are present and strategy δ is used by all (including future arrivals).

Theorem 1. Suppose that $\lambda < \mu$ and $R < C/(\mu - \lambda)$. Let p_i, $i \geqslant 0$, be the probability that a customer who observes i customers in the system joins. Denote by $\underline{p}$ the corresponding vector. Let $W_i(\underline{p})$ be the expected waiting time for a customer who joins after observing i customers when all others behave according to $\underline{p}$. Then, there exists a

[7] This value coincides with the expected time in random queues where the customer to whom service is granted is selected randomly among all those in the system at the time of service completion. The analysis here also suits this model.

[8] Footnote 6 applies here too, and reneging is forbidden here as part of the models' assumptions. For the case where reneging is allowed see [6].

[9] The model in [10] is more general. In particular, it assumes a $GI/M/s$ LCFS system without preemption. In the single server model that we assume, there is no qualitative difference between systems with and without preemption.

unique equilibrium in which $0 < p_i < 1$ for every $i \geqslant 0$. In this equilibrium, $p_i = p_e$ for every $i \geqslant 0$, where

$$p_e = \frac{1}{\lambda}\left(\mu - \frac{C}{R}\right). \tag{3}$$

Proof. Since $p_i > 0$ for every $i \geqslant 0$, the states compose a single communicating class and hence either they are all recurrent or they are all transient. The theorem's condition implies that $W_i(\underline{p})$, $i \geqslant 0$, are finite. Therefore all states are recurrent. For $0 < p_i < 1$ in equilibrium, customers must be indifferent between joining and balking in all states. Therefore, $R - CW_i(\underline{p}) = 0$ for every $i \geqslant 0$.

Consider a customer who joins after observing state $i - 1$. If the next event is an arrival then this arrival will preempt his service. He will then have to wait an expected time of $W_i(\underline{p})$ until his service is resumed. Therefore,

$$W_{i-1}(\underline{p}) = \frac{1}{p_i\lambda + \mu} + \frac{p_i\lambda}{p_i\lambda + \mu}\big(W_i(\underline{p}) + W_{i-1}(\underline{p})\big).$$

Substituting $W_i(\underline{p}) = W_{i-1}(\underline{p}) = R/C$ we get that for every $i = 0, 1, \dots$

$$\frac{R}{C} = \frac{1}{\mu - p_i\lambda},$$

and therefore $p_i = p_e$ where p_e satisfies (3). □

Note that since under the assumption of theorem 1 a single communicating class is formed, the resulting equilibrium is also a SPE.

We now investigate pure SPE's. If $\lambda < \mu$ and all arriving customers join, the expected waiting time (including the time in service) of a customer is $1/(\mu - \lambda)$. If also $R \geqslant C/(\mu - \lambda)$, then joining is a best response in any state and against any strategy. In other words, the strategy of always joining is a dominant strategy and no other equilibrium exists.

Assume now that either $\lambda \geqslant \mu$ or $\lambda < \mu$ and $R < C/(\mu - \lambda)$. For $m \geqslant 1$, let $B(m)$ be the expected waiting time for a customer who arrives to an empty LCFS-PR queue when all use the strategy of joining if and only if the number of customers in the system is at most $m - 1$ (and balk otherwise). It is clear that $B(m)$ is monotone increasing with m.[10] This value is also equal to the expected waiting time of a customer who joins when there are k customers in the system and the strategy used by all prescribes joining at states $k+1, \dots, k+m-1$ and balking at state $k+m$ (it is immaterial what the strategy prescribes for other states). Let $m^* = \max_{m\geqslant 1}\{m \mid R \geqslant CB(m)\}$. Note that m^* exists by our assumption that either $\lambda \geqslant \mu$ or $\lambda < \mu$ and $R < C/(\mu - \lambda)$. A pure strategy can be written as a string of 0's and 1's, where 0 denotes balking and 1 denoted joining.

[10] In fact, $B(m)$ coincides with the expected length of a busy period in an $M/M/1/m$ model. This value is $\sum_{i=0}^{m-1} \rho^i/\mu$ (see, for example, [7]).

The following two conditions are necessary for a pure strategy to be a SPE:

- Each 1 is followed by at most $m^* - 1$ 1's; otherwise, this 1 is suboptimal.
- Each 0 is followed by at least m^* 1's; otherwise, this 0 should be replaced by 1 for a better action.

Combining these conditions, each 0 has to be followed by *exactly* m^* 1's which are followed by a 0. This condition is also sufficient for a SPE. Thus, there are exactly $m^* + 1$ pure SPE solutions. All of them are based on repeating indefinitely identical "blocks" of length $(m^* + 1)$, each having a 0 followed by m^* 1's. They differ by where the number of 1's before the first 0. In order to obey the first of the two conditions, the 0 initiating the first block, should appear in any of the first $m^* + 1$ entries.

If ρ is sufficiently large, $m^* = 1$ and hence two pure SPE solutions exists. The first is $(0, 1, 0, 1, \dots)$, prescribing to join even queue sizes and balk at odd queue sizes. Under this strategy the server is always idle! The other pure SPE is $(1, 0, 1, 0, \dots)$.

We describe now some illustrative pure equilibria that are not subgame perfect. Suppose the arrival rate is large. Then, "join unless $s = 0$" is an equilibrium, but of course not a SPE. In fact, under this strategy nobody joins (assuming the process initiates with an empty system) and $s = 0$ is the only recurrent state. Likewise, the strategy "join unless $s = 1$" is an equilibrium but not a SPE. Note that under this strategy at most one customer is present in the system at any given time, and no preemption takes place. Indeed, the only recurrent class here is $\{0, 1\}$. Lastly, the strategy "join unless $s = 2$" is not an equilibrium because when this strategy is adopted by all, joining when $s = 0$ (which is a recurrent state) is not an optimal response due to the high arrival rate. If the arrival rate λ is reduced to a moderate value,[11] this strategy becomes an equilibrium.

Finite buffer

Assume that for some number, $N \geqslant 1$, a customer who observes upon arrival N customers in the system (including the one in service) must balk. This is equivalent to assuming that the waiting room (including the service position) cannot accommodate more than N customers, and new arrivals are rejected. A more general version of this model was treated in [10] and a single equilibrium, which turns out to be an SPE, was computed. We give here a more complete analysis of the pure SPE's in this case.

Here, $S = \{0, 1, \dots, N - 1\}$ and $R_s(0, \delta) = 0$ for $s \in S$. Let δ be a pure strategy. For a state s, let state $i(s, \delta)$ be the smallest state larger than or equal to s for which δ prescribes balking. If such a state does not exist, set $i(s, \delta) = N$. The expected waiting time of a customer who joins after observing state s is $B(i(s+1, \delta) - s + 1)$. Therefore,

$$R_s(1, \delta) = R - CB\big(i(s + 1, \delta) - s + 1\big), \quad 0 \leqslant s \leqslant N - 1.$$

If $m^* \geqslant N$, then "always join" is the unique equilibrium. Moreover, under this strategy all states are recurrent and therefore this is also a SPE.

[11] For the exact condition see [10, p. 492].

Suppose now that $m^* < N$. For a pure strategy δ, denote by $n(\delta)$ the first entry where 0 appears. The following two conditions are necessary and sufficient for δ to be an equilibrium:

- $n(\delta) \leqslant m^*$.
- 1 appears in positions $n(\delta) + 1, \ldots, \min\{N, n(\delta) + m^*\}$.

Under a strategy which obeys the two conditions, states $\{0, \ldots, n(\delta)\}$ are the only recurrent states and indeed an optimal action is prescribed for them. Which among them are SPE? The answer is simple: Insert 1's at the m^* positions starting from $N - 1$ and going down, then insert a 0, then an additional m^* of 1's, etc.

2.3. *Example 4: Two priority classes*

The following model was considered by Adiri and Yechiali [1] and Hassin and Haviv [7]. It is similar to the model in example 1 but balking is not allowed and customers have the option of purchasing priority at the cost of θ. They make their decisions after observing the state (i, j) where i and j denote the number of ordinary and priority customers, respectively, in the system. A priority customer who observes upon his arrival an ordinary customer in service, commences service immediately while preempting the customer in service. A preempted service is resumed at the point where it was interrupted. Customers belonging to the same priority class are served in a FCFS order. For stability we assume that $\lambda < \mu$.[12]

Denote the action of not purchasing priority by 0 and that of purchasing priority by 1. Then, for any $i, j \geqslant 0$, $R_{(i,j)}(1, \delta) = C(j+1)/\mu$ and $R_{(i,j)}(0, \delta) = CW(i+1, j, \delta)$ where $W(i, j, \delta)$ is the expected time in the system for an ordinary customer who is in position i in his class and when additional j priority customers are present, given that all use strategy δ. The expected waiting time for an ordinary customer depends on the strategies selected by future arrivals, and hence it is not a trivial function of δ.

In [7] it is proved that there exist integers $n_1 \leqslant n_2$ such that for any integer n, $n_1 \leqslant n \leqslant n_2$, the following is a pure equilibrium strategy: $\delta(0|(i, j)) = 1$ if $i < n$, otherwise, $\delta(1|(i, j)) = 1$. Under δ the recurrent states are $(i, 0)$ for $0 \leqslant i \leqslant n$, and (n, j) for $j \geqslant 0$. Assuming $(0, 0)$ to be the initial state, then priority is not purchased under δ as long as the total (and hence, low priority) number of customers in the system does not reach the threshold n. Priority is purchased when this number is grater than $n - 1$.

We observe that other pure equilibria exist. For example, for a strategy δ' which is similar to δ except for that $\delta'(1|i, j) = 0$ for $i \geqslant n + 1$, is an equilibrium. The sets of recurrent states under δ and under δ' coincide, and the two strategies disagree only in transient states (see remark 1).

[12] Since balking is not allowed the value of service plays no role. For this reason, in this example and the next one, we minimize the system's costs rather than maximizing welfare, as we have done in the previous examples.

The next question is whether δ is a SPE. It is true that δ prescribes optimal responses for all transient states (i, j) with $i > n$ but this is not necessarily so with respect to the transient states (i, j) where $i < n$. Consider the transient state under δ, $(n-1, j)$ with a large value of j. For this state, δ prescribes not to purchase priority. Yet, since any priority customer present adds (under δ) a busy period of waiting time to each ordinary customer, for a large j, the only optimal response is to purchase priority.

A SPE strategy is based on an integer n as described above, coupled with additional n thresholds denoted by $j(n, n-1) < j(n, n-2) < \cdots < j(n, 0)$ such that for states (i, j) with $i \leqslant n-1$ and $j \geqslant j(n, i)$ it prescribes purchasing priority. Otherwise, the strategy will be as δ above.

We next show how to compute $j(n, n-1)$. Finding $j(n, n-2)$, $j(n, n-3)$, $\ldots, j(n, 0)$ is more complicated and not accomplished here. We learn from [7] that a strategy δ with a threshold n as defined above is an equilibrium if and only if

$$\frac{C}{\mu} + \theta - CB \leqslant CW(n) \leqslant \frac{C}{\mu} + \theta,$$

where $W(n)$ is the expected waiting time of an ordinary customer who is the last among n (including himself) ordinary customers when no priority customer is present, and where $B = 1/(\mu - \lambda)$ is the expected length of a busy period (given that all customers follow δ). Consider now a customer who observes a transient state $(n-1, j)$. If he purchases priority his expected cost is $\theta + C(j+1)/\mu$. Otherwise, it is $jCB + W(n)$. Thus, for $j \geqslant j(n, n-1)$ where

$$j(n, n-1) = \min_{j \geqslant 1} \left\{ \theta + C\frac{j+1}{\mu} \leqslant jCB + CW(n) \right\},$$

a SPE prescribes purchasing priority.

2.4. Example 5: Choosing between batch servers

Suppose there are two means of transportation. The first is a bus of infinite capacity for which the mean waiting time is five minutes. The second is a seven-seater shuttle which departs as soon as the seventh commuter joins. There is no shortage of shuttles and as soon as a shuttle departs another one arrives. The commuters' arrival process is Poisson with a rate of λ per minute. Each commuter, after observing the occupancy level of the shuttle, selects one of the means of transportation.

$S = \{0, 1, \ldots, 6\}$ is the occupancy level of the shuttle. Action 0 is to select the bus and action 1 is to join the shuttle. For simplicity, we deal only with pure strategies. Let $n(\delta)$ be the largest state for which δ prescribes taking the bus: $n(\delta) = \max_{s \in S}\{(0|s) = 1\}$. If δ always prescribes choosing the shuttle, set $n(\delta) = -1$. Then $R_s(0, \delta) = 5$ for all $s \in S$ and all $\delta \in \Delta$, and

$$R_s(1, \delta) = \begin{cases} \infty, & 0 \leqslant s < n(\delta), \\ \dfrac{6-s}{\lambda}, & 6 \geqslant s \geqslant n(\delta). \end{cases}$$

If $n(\delta) \geqslant 0$ then the only recurrent states are those with $\delta(0|s) = 1$. These are absorbing states and thus each of them constitutes a recurrent class of its own. When $n(\delta) = -1$, i.e., $\delta(1|s) = 1$ for all $s \in S$, S constitutes a single recurrent class.

Consider the strategy under which all use the shuttle. The expected waiting time for one who observes an empty shuttle is $6/\lambda$. Thus, this strategy is an equilibrium if and only if $\lambda > 6/5$. Moreover, since all states are recurrent, it is a SPE as well. It is the only equilibrium when $\lambda > 6/5$: indeed when all others use some other strategy δ, selecting the bus at the recurrent state $n(\delta)$ is a suboptimal response.

As for other strategies, note that δ is not an equilibrium if $n(\delta) = 6$ as it prescribes taking the bus when facing the recurrent state 6 (and not the shuttle which will depart immediately). If $n(\delta) \leqslant 5$, check if $[6 - n(\delta)]/\lambda \geqslant 5$. If this is the case then δ is an equilibrium. Otherwise, it is not. Once δ is found to be an equilibrium, the next issue is whether it is also a SPE. Two conditions are to be met here. First, it is required that $\delta(0|s) = 1$ for all $s \leqslant n(\delta)$ (so optimal responses are prescribed for all states $s \leqslant n(\delta)$ which are recurrent), and second, that $\delta(1|s) = 1$ for all states for which if all future arrivals take the shuttle, the optimal response is also to take the shuttle. These states which are with $s > n(\delta)$, are transient.

It is clear that only one pure SPE exists. Specifically, we already considered the case $\lambda > 6/5$ and concluded that the unique pure SPE is that all use the shuttle. Otherwise, a SPE requires taking the shuttle when, and only when, the number of waiting commuters there is i with $(6 - i)/\lambda < 5$, i.e., $i > 6 - 5\lambda$. For example, when $\lambda < 1/5$ under the SPE one selects the shuttle if and only if it will depart immediately.

3. Concluding remarks

The concept of subgame perfection has been introduced by game theorists to refine the solution concept of Nash equilibrium. The idea is to rule out some equilibria and end up with the more “rational” ones.

For a given equilibrium, there are some (information) states that are out of the equilibrium path and hence an irrational prescription of actions for such states does not necessarily lead to a violation of the equilibrium requirements. These states may be considered as reachable only due to an error in the execution of the equilibrium profile (known as the “trembling hand effect”). For subgame perfection, it is required that rational (i.e., optimal) actions are prescribed also for these states. This paper focused on this aspect of subgame perfection and its applications to observable queues.

A second issue is that sometimes equilibrium paths are themselves irrational and, in fact, are based on incredible threats. For example, in the *Entry Game*, one player threatens to flood the market if a second player enters it. However, given that the second player has chosen to enter (in spite of the threat), flooding the market is a sub-optimal response for the first player and hence the threat strategy is not a SPE. There are interesting queueing games where the concept of threats is meaningful. Investigating such models is an interesting problem that this paper leaves for future research. We conclude however with some examples.

Consider the processor sharing model of section 2.1 with homogeneous customers, exponential service requirements, and given time values and benefits from service. An arriving customer decides whether to join or balk, after observing the queue. If the queue is sufficiently long, then joining of the new arrival may be associated with a negative expected net benefit to him, as well as to those already in the queue. Upon his arrival, the customer has the option of balking. But, by the memoryless property, once he joins he cannot be distinguished from the others. This model was analyzed by Altman and Shimkin [2] who computed an equilibrium threshold strategy. Some details were given above in section 2.1.

We now add the option of reneging: at any time, a customer in the queue is allowed to leave at no additional cost or benefit. This assumption puts the new arrival in the same position as any other customer who is already in the queue. This model was investigated by Assaf and Haviv [3] under the assumption that the only information customers possess while they are in the system is the queue size. In [3] it is proved that no SPE exists (although the concept of SPE was not discussed there) in this model, and, consequently, a weaker concept of equilibrium was defined and solved.

We will now show that the model has an equilibrium. Consider the threshold, n, of Altman and Shimkin. The following is an equilibrium strategy: join if and only if the number of customers you observe upon arrival is less then n; otherwise balk; never renege.

The fact that a customer can renege works against him, since it encourages new arrivals to enter while hoping that he will renege. Once a customer announces that he will never renege, others may not join and he will be better off. Of course this isn't a SPE since if they join, staying in the system is suboptimal. Thus, when we forbid equilibrium strategies that are not SPE we actually move from the model of Altman and Shimkin to the model of Assaf and Haviv.

In another example, suppose that customers arrive to a processor sharing system, observe the queue and decide whether to join. A customer who joins also selects his service requirement. A pricing scheme that induces the socially optimal behavior in equilibrium in this model is given by Ha in [5]. Assume that a longer service gives the customer an increased value, but the value function is concave. When the number of customers in the queue increases, the rate of service each of them receives decreases, and this may induce those who entered earlier to renege. Suppose that an arriving customer has full information on the time already spent in service by the customers in the system.[13] Based on this information, the new arrival can compute the consequences that his joining has on the others, in particular their expected times of leaving the queue. Using this information he may join or balk. The customers in the system are interested in deterring the new arrival from joining. One way of doing so is by threatening that if he joins they will stay for long time (possibly as long as he is in the queue). Such a strategy may define an equilibrium but it isn't a SPE.

[13] The model is of equal interest if this information is not available.

References

[1] I. Adiri and U. Yechiali, Optimal priority purchasing and pricing decisions in non-monopoly and monopoly queues, Operations Research 22 (1974) 1051–1066.

[2] E. Altman and N. Shimkin, Individual equilibrium and learning in processor sharing systems, Operations Research 46 (1998) 776–784.

[3] D. Assaf and M. Haviv, Reneging from time sharing and random queues, Mathematics of Operations Research 15 (1990) 129–138.

[4] K.R. Balachandran, Purchasing priorities in queues, Management Science 18 (1972) 319–326.

[5] A.Y. Ha, Optimal pricing that coordinate queues with customer-chosen service requirements, Management Science 47 (2001) 915–930.

[6] R. Hassin, On the optimality of first come last served queues, Econometrica 53 (1985) 201–202.

[7] R. Hassin and M. Haviv, Equilibrium threshold strategies: The case of queues with priorities, Operations Research 45 (1997) 966–973.

[8] R. Hassin and M. Haviv, *To Queue or Not to Queue: Equilibrium Behavior in Queues* (Kluwer, 2002) (forthcoming).

[9] P. Naor, The regulation of queue size by levying tolls, Econometrica 37 (1969) 15–24.

[10] B. Tilt and K.R. Balachandran, Stable and superstable customer policies with balking and priority options, European Journal of Operational Research 3 (1979) 485–498.