

Minimizing the worst slowdown: off-line and on-line

Hervé Moulin
Rice University

April 2005

Abstract

Minimizing the slowdown (expected sojourn time divided by job size) is a key concern of fairness in scheduling and queuing problems where job sizes are very heterogeneous. We look for protocols (service disciplines) capping the worst slowdown (called here *liability*) a job may face no matter how large (or small) the other jobs are.

In the **scheduling** problem (all jobs released at the same time), allowing the server to randomize the order of service cuts almost in half the liability profiles feasible under deterministic protocols. The same statement holds if cash transfers are feasible and users have linear waiting costs.

In a **queuing** problem (release times of jobs are arbitrary), we can construct a deterministic **on-line** (non anticipative) protocol guaranteeing the liability θ_r to job i , where r is the number of jobs in the queue when i was released, *if and only if* $\sum_1^\infty \frac{1}{\theta_r} \leq 1$. When the arrival of new jobs is Poisson with rate λ , the liability of a job of size x is no smaller than its slowdown when all other jobs are of the same size, namely $\frac{1}{1-\lambda \cdot x}$. We conjecture that this liability is feasible on-line, and identify a probabilistic protocol achieving the liability $\frac{1.45}{1-\lambda \cdot x}$.

Key words : scheduling, queuing, slowdown, probabilistic protocols

Acknowledgement 1 *Thanks to Eric Friedman, Jay Sethuraman and seminar participants at Columbia University, the University of Guanajuato, Mexico, and the Roy-Malinvaud seminar in Paris.*

1 Individual Guarantees

Several agents share resources according to a mechanical rule of which the input is the profile of individual characteristics. The *guarantee* of a particular agent is the smallest welfare/utility level she will reach, under the worst possible configuration of other agents' characteristics. This level only depends upon this agent's own characteristics, the resources to be shared, and the number of other agents.

To a participant with no information about the other agents with whom resources are shared, the guarantee is a simple measure of her downside risk, influencing both her willingness to participate in the mechanism and her perception of its fairness. Therefore an important design criteria is to improve guarantees as much as permitted by the nature of the resources being allocated. This criteria is as old as the fair division literature, and inspires for instance the familiar "I-Divide-You-Choose" mechanism¹. More discussion in a variety of micro-economic allocation problems can be found in [20], [22], [6], [13], [14], [15].

Here we apply the idea of maximizing individual guarantees to general scheduling and queuing problems. A single server is the shared resource. An agent's characteristics are the size (processing time) of the job he submits to the server and its release date (in the case of queuing). The expected sojourn time, from release of the job until completion by the server, is the agent's disutility.

A central debate on the management of queues bears on congestion control in the presence of "ill-behaved sources" ([5]), namely queues where the service time may vary wildly across different users. Two conflicting normative goals inspire the discussion: to minimize the sum of individual sojourn times, or to equalize the *slowdown* (sojourn time divided by service time) of the different users. The former is the classic utilitarian minimization of the sum of individual disutilities, the latter applies an egalitarian concern to slowdowns, namely excess wait *per unit of job*, thus viewing as fair that a job three times larger stays in the system three times longer. See [5] [3], [1], [8], [23], [24] and references therein.

The benchmark protocols *Shortest Job First* (SJF) ², and *Shortest Re-*

¹If utilities are additive over the pieces of a cake, Divide and Choose guarantees to each agent a piece worth at least half of the entire cake.

²in a *scheduling* problem, where all jobs are released at the same time.

maintaining Job First (SRJF)³ are utilitarian optimal, but it is often argued that they are too harsh on long jobs: the resulting slowdowns favors lexicographically the smaller jobs. See [1], [2], [8], [10]. An exciting stream of recent research discusses service protocols achieving a reasonable compromise between the two conflicting goals. In an M/G/1 queue the familiar *Processor Sharing* protocol (PS) (serving all active jobs at the same rate) equalizes (expected) slowdown across all users. Not only is PS very far from utilitarian optimal, it also uses the server inefficiently because partially completed jobs are useless. It is not hard however, to design an efficient protocol Pareto superior to PS: this is the *Fair Sojourn Processing* (FSP) introduced in [9] (see also [5]). FSP achieves a nearly optimal total sojourn time, while guaranteeing to every user a smaller slowdown than PS ([8], [9], [23]).⁴

In this paper we take a different route, and compare various service disciplines by the guarantees they offer to the users. That is, we focus on the worst slowdown - we call it the *liability*- that a given user may experience, where the minimum is taken over all possible distributions of job sizes and release dates for other users. We submit that in many real-life queues involving heterogenous users, such as the internet, ignorance of other users' characteristics is the norm rather than the exception. Many human queues are subject to unpredictable bursts and lapses, and the service time may differ widely across users. Insisting that the service protocol minimizes individual liabilities is the simplest way to protect individual users against the unknown, potentially very large, heterogeneity of individual demands.

We are looking for protocols guaranteeing a bounded liability, under the most parsimonious informational assumptions.

In a scheduling problem (sections 3,4) individual liability will only depend upon the number of other users, or their name. We compute the minimal feasible liability profiles under a deterministic protocol, then when randomization is allowed (and users care about the expected sojourn time), and finally when cash transfers (but not randomization) are feasible and users have linear waiting costs. We find that randomization (Theorem 1), *or* cash transfers (Theorem 2), cut nearly in half the liability profiles feasible under deterministic protocols.

In a queuing problem (sections 5,6,7), we consider *on-line* (non anticipa-

³in a *queuing* problem, where release dates vary.

⁴These features of FSP hold in any queue, no matter how the successive jobs are released. Yet when the arrival of new jobs is not Poisson, FSP may not result in an egalitarian, or nearly egalitarian, profile of slowdowns.

tive) protocols of two different types. If nothing at all (number, sizes and release dates) is known about future jobs, we can still offer to any job a finite liability that only depends upon the number of live (unfinished) jobs at release time. We construct in Theorem 3 a deterministic protocol to that effect. An alternative, very common assumption throughout the entire queuing literature, is that the arrival of future jobs follows a Poisson process with known rate. In the steady-state of the arrival process, the liability only depends upon own job size and the arrival rate of jobs. We compute a natural lower bound for any feasible liability, conjecture that this lower bound is feasible, and offer a protocol of which the liability is at most 45% larger: Theorem 4.

2 Overview of the results

We start with a scheduling problem involving a set N of users, and illustrate the concept of liability for two benchmark protocols. Consider the *Random Order* (RO) protocol, selecting all service orderings with equal probability. If the service time of job $i, i \in N$, is x_i , its expected sojourn is $x_i + \frac{1}{2} \sum_{j \neq i} x_j$. As x_j can be arbitrarily larger than x_i , the slowdown of job i is unbounded (its liability is infinite). By contrast *Shortest Job First* (SJF) offers the liability $n = |N|$ to every user: the worst case is when all jobs $j, j \neq i$, are barely shorter than x_i , implying that i is served last and her slowdown is n . If it is desirable to treat users unequally, we implement similarly the profile of liabilities $(\theta_i, i \in N)$ by the *Shortest Deadline First* (SDF) protocol, serving the jobs in the order of their respective "deadlines" $\theta_i x_i$. This is feasible if and only if $\sum_N \frac{1}{\theta_i} \leq 1$: Proposition 1. These bounds cannot be improved by any deterministic protocol, but they can be cut nearly in half for a *probabilistic* protocol, where the server randomizes the service ordering.

To see why randomization is useful here, consider n jobs of size 1. If one of these jobs is deterministically scheduled last, its slowdown is n . But if all orderings are equally probable, the common slowdown is $\frac{n+1}{2}$. More generally, the profile of liabilities $(\theta_i, i \in N)$ is probabilistically feasible if and only if $\sum_N \frac{1}{2\theta_i - 1} \leq 1$: Theorem 1. We offer a couple of simple protocols to implement these liabilities: their key property is that the expected delay imposed by job j on job i (i.e., the probability that j is scheduled before i , multiplied by x_j) never exceeds $\frac{1}{2} \frac{2\theta_i - 1}{2\theta_j - 1} x_i$. See Proposition 2.

Randomization is the easiest way to restore fairness when jobs must be

processed whole. An alternative device is cash transfers, of which the implementation is more difficult. It is clearly impractical in very large problems such as the Internet; an additional hurdle is to elicit individual tradeoffs between delay and cash compensation. Here we make the familiar simplifying assumption that each users' waiting cost per unit of time is constant and known to the server⁵. The net cost of a user is the sum of the cost of his service time and a cash transfer, and his slowdown is the ratio of net cost to the "stand alone" cost of service (when he has the server all to himself). The model with cash transfers and linear waiting costs is profoundly different from the probabilistic one. In the latter all probability distributions on the service ordering yield efficient (Pareto optimal) profiles of expected service time, whereas in the former efficiency of the profile of net costs essentially determines the service ordering. Thus it comes as a surprise that the feasible liability profiles $(\theta_i, i \in N)$ in the two models are precisely the same: Theorem 2.

We turn to the queuing problem. The case of an *off-line* service protocol, where the server knows at time 0 the size and release date of all present and future jobs, is a simple variant of the scheduling model, to which our Theorems 1 and 2 extend word for word. But in most real-life queues, the users and the server have very little information about future jobs, and a feasible protocol must be *on-line*, i.e., non-anticipative. We have something to say about the smallest feasible liability in two important special cases.

Assume first that the server - and the users- know absolutely nothing (number, size or release date) about future jobs. In particular the flow and size of future jobs may create unlimited congestion that we are unable to predict today. The challenge is to construct a protocol based on the state of the queue at release time, and guaranteeing a bounded expected slowdown.

We note first that none of the familiar on-line protocols discussed in the literature achieves does the job, i.e., they all yield an infinite liability. Start with SRJF, that offers no protection against the release of an arbitrarily large number of jobs shorter than my own; this is the argument, mentioned above, that while helping small jobs, SRJF unfairly penalizes large jobs⁶. The same observation applies to PS (and to its FSP improvement). For instance in the M/G/1 queue, the common slowdown under PS is unbounded if the queue

⁵See [7], [21], [11], [4], for a discussion of fair and incentive-compatible cash compensations in that context

⁶See [1], [3], [24], and the recent work estimating the resulting slowdowns under several assumptions on the sizes and release dates of new jobs.

is unstable⁷. Under *First Come First Serve* (FCFS), a newly-released job must wait until all live jobs are completed, no matter how large these could be. Finally under *Last Come First Serve*(LCFS), a very long job released just after me, or a burst of relatively small jobs, will increase my service time without bounds, just like under SRJF.

Using a weighted version of the FSP protocol ([8],[9]), we can nevertheless cap the liability of any user as a function only of the size r of the queue (the number of live jobs) when his job is released. If $\{\theta_1, \theta_2, \dots\}$ is a sequence of positive numbers such that $\sum_1^\infty \frac{1}{\theta_r} \leq 1$, our protocol guarantees the liability θ_r to any user who is released in a queue of size r . Conversely, this is only possible when the above inequality holds: Theorem 3. Taking r as a proxy for the congestion in the queue, we conclude that the liability must increase more than linearly in r , in sharp contrast with the scheduling context.

Finally we consider the much discussed context (often called the M/G/1 queue; see [25]), where the release of new jobs follows a Poisson process with rate λ known to the server (and users). We look for a finite liability that only depends upon the arrival rate λ and own job size x (we normalize the service rate to 1). As jobs can be arbitrarily large, the queue may well be unstable. If all jobs in the queue are of identical size x , the queue is stable if and only if $\lambda x < 1$, and in this case the expected slowdown of each job under a fair protocol⁸ is $\frac{1}{1-\lambda x}$.

Two observations follow. To a job larger than $\frac{1}{\lambda}$, we cannot offer a finite liability (depending only upon λ and x). For a job smaller than $\frac{1}{\lambda}$, the liability cannot be less than $\frac{1}{1-\lambda x}$. I conjecture that there exists a probabilistic on-line protocol providing precisely this liability to every job $x, 0 \leq x < \frac{1}{\lambda}$. Theorem 4 describes one such protocol achieving the liability $\frac{1.45}{1-\lambda x}$ for all $x, 0 \leq x < \frac{1}{\lambda}$. It decides randomly where a given job must be inserted in the existing queue, and the larger the job, the further back it is likely to be pushed.

⁷With an arrival rate λ and mean job size \bar{x} , stability requires $\lambda \bar{x} < 1$, and in this case the common slowdown is $\frac{1}{1-\lambda \bar{x}}$ (e.g., [25]).

⁸Such as FCFS, LCFS, RO, or any strong and work conserving discipline treating two jobs equally if they face the same queue, see [25].

3 Probabilistic scheduling

There is a single server, with service rate normalized to 1, and the processing time of a job is deterministic⁹. A *scheduling problem* is a pair (N, x) , where N is a finite set of users, and $x = (x_i, i \in N)$ is a profile of (strictly) positive job sizes. Users care only about their sojourn time, namely the date of completion of their own job (a partially completed job is useless). Therefore preemption is inefficient, jobs must be served whole.

A *probabilistic* (resp. *deterministic*) *protocol* π associates to every problem (N, x) a random ordering $\sigma = \pi(N, x)$ of N , namely a probability distribution on the set of orderings of N (resp. an ordering of N). Denoting $P(i, \sigma)$ the (random) set of agents preceding i in σ (including i itself), the expected sojourn time of job i is thus $y_i = E_\sigma[\sum_{P(i, \sigma)} x_j]$.

Given a protocol, and a problem (N, x) , the expected *slowdown* of job i is $s_i(N, x) = \frac{y_i}{x_i}$. Given N , a profile $\theta = (\theta_i, i \in N)$ is a *feasible liability profile* if there exists a protocol π such that

$$s_i(N, x) \leq \theta_i \text{ for all } x \gg 0, \text{ and all } i \in N \quad (1)$$

A feasible liability profile is *minimal* if for any different feasible liability profile θ' , $\theta'_i > \theta_i$ for some $i \in N$.

Our first result characterizes the feasible liability profiles when the server is restricted to deterministic protocols.

Proposition 1

Given the set N of users, the vector θ is a minimal feasible liability profile for a deterministic protocol, if and only if $\sum_N \frac{1}{\theta_i} = 1$. In this case it is implemented by the Earliest Deadline First protocol, serving job i before job j only if $\theta_i x_i \leq \theta_j x_j$.

Proof

If θ is a feasible liability profile, choose $x_i = \frac{1}{\theta_i}$ for all i and apply (1) to the agent i ranked last: this gives $\sum_N \frac{1}{\theta_i} \leq 1$. Conversely, assume $\sum_N \frac{1}{\theta_i} = 1$, and apply the earliest deadline first protocol, breaking ties arbitrarily. For the problem (N, x) , this protocol will select an ordering $\{1, \dots, n\}$ of N such that $\theta_1 x_1 \leq \theta_2 x_2 \leq \dots \leq \theta_n x_n$. For any k , we have then

⁹All our results are preserved if the processing time of a job is random, provided the time it takes to process different jobs are stochastically independent.

$$\sum_1^k x_j \leq \sum_1^k \frac{\theta_k x_k}{\theta_j} \leq \theta_k x_k$$

thus our protocol implements θ . ■

Allowing probabilistic protocols reduces by nearly 100% the minimal liability profiles.

Theorem 1

Given the set N of users, the vector θ is a minimal feasible liability profile for a probabilistic protocol, if and only if $\sum_N \frac{1}{2\theta_i - 1} = 1$.

Proof

Step 1: a preliminary result. Given a problem (N, x) , write $F(N, x)$ for the set of feasible profiles of expected sojourn times, namely

$$y \in F(N, x) \Leftrightarrow \text{for some random ordering } \sigma, y_i = E_\sigma[\sum_{P(i, \sigma)} x_j] \text{ for all } i$$

Define for all $x \in \mathbb{R}_+^N$ and all $S \subseteq N$, the function $v(S, x) = \sum_S x_i^2 + \sum_{S(2)} x_i \cdot x_j$, where $S(2)$ is the set (with cardinality $\frac{|S| \cdot (|S| - 1)}{2}$) of non ordered pairs from S . Note that v is supermodular with respect to S . The following result is proven in [17] and [16].

$$y \in F(N, x) \Leftrightarrow \left\{ \sum_N x_i \cdot y_i = v(N, x) \text{ and } \sum_S x_i \cdot y_i \geq v(S, x) \text{ for all } S \subseteq N \right\}$$

Step 2: only if statement. Let θ be feasible at N . For any x there exists $y \in F(N, x)$ such that $y_i \leq \theta_i x_i$, therefore by step 1

$$\sum_N \theta_i \cdot x_i^2 \geq v(N, x) \Leftrightarrow \sum_N (2\theta_i - 1) \cdot x_i^2 \geq \left(\sum_N x_i \right)^2 \quad (2)$$

For $x_i = \frac{1}{2\theta_i - 1}$, this inequality reduces to $\sum_N \frac{1}{2\theta_i - 1} \leq 1$.

Step 3: if statement. Fix θ such that $\sum_N \frac{1}{2\theta_i - 1} \leq 1$, and a problem (N, x) . Because the function v is supermodular, the core of the game $(N, v(\cdot, x))$ is "large" (see [18]). That is, for any fixed $z \in \mathbb{R}_+^N$ there exists $y \in F(N, x)$ such that $y_i \leq z_i$ for all i if and only if $\sum_S z_i \cdot x_i \geq v(S, x)$ for all $S \subseteq N$. Choosing $z_i = \theta_i \cdot x_i$, we find that the slowdown $(\theta_i \cdot x_i)$ is feasible at (N, x) if (2) holds for all subsets S of N , including N itself. Choose S , define $u_i = \sqrt{2\theta_i - 1} \cdot x_i$, $w_i = \frac{1}{\sqrt{2\theta_i - 1}}$, for all $i \in S$, and apply Schwartz's inequality to u and w :

$$(u \cdot w)^2 = \left(\sum_S x_i\right)^2 \leq \|u\|^2 \cdot \|w\|^2 = \left(\sum_S (2\theta_i - 1) \cdot x_i^2\right) \cdot \left(\sum_S \frac{1}{2\theta_i - 1}\right)$$

Thus our choice of θ guarantees the desired inequality (2). ■

An important special case of the two results above is that of an anonymous liability, $\theta_i = \bar{\theta}$ for all i . For deterministic protocols, the minimal feasible liability is $\bar{\theta} = n$, and is implemented by SJF. For probabilistic protocols, it is $\bar{\theta} = \frac{n+1}{2}$, and is implemented by the anonymous version of the parametric protocols to which we now turn.

Fix the set N of users and choose for each i and job size $x_i > 0$ a cumulative distribution function $F_{i,x_i}(z)$ on $[0, +\infty[$. That is, F_{i,x_i} is any non negative, non decreasing and right-continuous function on $[0, +\infty[$ such that $\lim_{\infty} F_{i,x_i}(z) = 1$. Given a problem (N, x) , the corresponding *parametric protocol* draws for each user i a random variable Z_i according to F_{i,x_i} , and these draws are stochastically independent. It then serves the jobs in the order of the realizations $Z_i(\omega)$: job j is served before job i if $Z_j(\omega) < Z_i(\omega)$, ties being broken by a fair coin. This rich family of protocols, which includes in particular SJF and RO, was introduced in [16], and studied there from the point of view of the strategic maneuvers of splitting and merging jobs.

Proposition 2

Given N , a feasible liability profile θ as in Theorem 1 is implemented by the following two parametric protocols:

$$\begin{aligned} \text{weighted quadratic:} \quad & F_{i,x_i}(z) = \min\{z^{(2\theta_i-1)^2 x_i^2}, 1\} \text{ for } 0 \leq z < \infty, \\ \text{weighted serial} \quad & : F_{i,x_i}(z) = \min\left\{\frac{z}{(2\theta_i-1)x_i}, 1\right\} \text{ for } 0 \leq z < \infty. \end{aligned}$$

Proof

Fix (N, x) and θ such that $\sum \frac{1}{2\theta_i-1} = 1$, and consider the weighted quadratic protocol. For any distinct agents i, j , write p_{ij} for the probability of the event $Z_j < Z_i$, which is easily computed as

$$p_{ij} = \frac{(2\theta_i - 1)^2 x_i^2}{(2\theta_i - 1)^2 x_i^2 + (2\theta_j - 1)^2 x_j^2} \tag{3}$$

Because the distributions F_{i,x_i} are atomless, p_{ij} is the probability that job j precedes job i . Note that it does not depend on the realizations of Z_k for $k \neq i, j$, and compute

$$\max_{x_j} p_{ij} \cdot x_j = \frac{2\theta_i - 1}{2\theta_j - 1} \cdot \frac{x_i}{2} \quad (4)$$

from which we get

$$y_i = x_i + \sum_{N \setminus i} p_{ij} \cdot x_j \leq \left(1 + \frac{2\theta_i - 1}{2}\right) \cdot \left(\sum_{N \setminus i} \frac{1}{2\theta_j - 1}\right) \cdot x_i = \theta_i \cdot x_i$$

A similar computation shows that equation (4) holds for the weighted serial protocol as well. ■

We stress that there are many more parametric protocols implementing a given liability profile θ . The key property is (4), namely the delay i expects to incur from j is at most the RHS term in (4). Besides the quadratic and serial c.d.f.s many other choices of F_{i,x_i} achieve this.

Equation (3) justifies our quadratic terminology: the probability that j precedes i is proportional to the square of the weighted job size. For an explanation of the serial terminology, see [16].

Remark 1

Proposition 1 generalizes easily to the case where N is countable, provided we assume that job sizes are integers $x_i = 1, 2, \dots$. Given a profile $\theta = (\theta_i, i \in N)$ such that $\sum_N \frac{1}{\theta_i} \leq 1$ and a problem (N, x) , we can always enumerate $N = \{i_1, i_2, \dots\}$ in such a way that $\theta_{i_1} x_{i_1} \leq \theta_{i_2} x_{i_2} \leq \dots$. Indeed the convergence of the series $\frac{1}{\theta_i}$, and $x_i \geq 1$ imply that for all a , the set of agents such that $\theta_i x_i \leq a$ is finite. Then the *Earliest Deadline First* protocol implements θ as above.

As for Theorem 1, it is preserved word for word if N is countable, and so is its proof.

4 Scheduling with cash transfers

The scheduling model in this section is very different, technically and in spirit, than the probabilistic model above. Randomization is not feasible, instead the server can perform cash transfers, provided they balance to zero. Each user has a linear disutility $\delta_i \cdot y_i - t_i$ over sojourn time and money, where δ_i is user i 's waiting cost per unit of time and t_i her cash transfer. Efficiency (Pareto optimality) in this context amounts to select an ordering of service

minimizing the sum of individual disutilities. Such orderings always serve job i before job j if $\frac{x_i}{\delta_i} < \frac{x_j}{\delta_j}$, and this property alone guarantees efficiency ([17]).

A scheduling *problem* is now a triple (N, x, δ) , where $x_i, \delta_i > 0$ for all i , and an *efficient protocol* associates to each problem (N, x, δ) a pair (σ, t) , where σ is an efficient ordering and $\sum_N t_i = 0$. The resulting net waiting cost w_i , and slowdown s_i of user i are

$$w_i = \delta_i \cdot \sum_{P(i, \sigma)} x_j - t_i, \text{ and } s_i(N, x, \delta) = \frac{w_i}{\delta_i x_i}$$

As in the previous model, $\theta = (\theta_i, i \in N)$ is a feasible liability profile if there exists a protocol such that $s_i(N, x, \delta) \leq \theta_i$ for all x and i .

Theorem 2

Given the set N of users, θ is a minimal feasible liability profile for a protocol with cash transfers, if and only if $\sum_N \frac{1}{2\theta_i - 1} = 1$.

Proof

Given a problem (N, x, δ) , the minimal total waiting cost (achieved by any efficient ordering) is easily computed as

$$W(x, \delta) = \min_{\sigma} \sum_N w_i = \sum_N \delta_i \cdot x_i + \sum_{N(2)} \min\{\delta_j x_i, \delta_i x_j\}$$

(recall that $N(2)$ is the set of non ordered pairs from N). Therefore θ is a feasible liability profile at N if and only if

$$\sum_N \theta_i \cdot \delta_i x_i \geq W(x, \delta) \text{ for all } x, \delta \gg 0 \tag{5}$$

Suppose (5) holds and pick $x = \delta$. As $W(x, x) = v(N, x)$ (defined in Step 1 of the proof of Theorem 1), we get $\sum_N \theta_i \cdot x_i^2 \geq v(N, x)$. Like in Step 2 of that proof, this implies $\sum_N \frac{1}{2\theta_i - 1} \leq 1$.

Conversely, we pick θ such that $\sum_N \frac{1}{2\theta_i - 1} \leq 1$ and prove (5). Change the variables δ_i to $\varepsilon_i = \frac{\delta_i}{x_i}$, so that we need to prove, for all $x, \varepsilon \gg 0$:

$$\sum_N \theta_i \cdot \varepsilon_i \cdot x_i^2 \geq \sum_N \varepsilon_i \cdot x_i^2 + \sum_{N(2)} \min\{\varepsilon_i, \varepsilon_j\} \cdot x_i x_j \tag{6}$$

When ε stays inside the cone $\{0 \leq \varepsilon_1 \leq \varepsilon_2 \leq \dots \leq \varepsilon_n\}$ of \mathbb{R}^N , inequality (6) is linear in ε , thus it is enough to prove it for the extreme directions of this

cone, namely $\varepsilon^1, \varepsilon^2, \dots, \varepsilon^n$, where $\varepsilon_i^k = 0$ if $i = 1, \dots, k$, $\varepsilon_i^k = 1$ if $i = k + 1, \dots, n$. Fix k and set $S = \{k + 1, \dots, n\}$. For ε^k , inequality (6) reduces to

$$\sum_S \theta_i \cdot x_i^2 \geq v(S, x)$$

which follows from $\sum_N \frac{1}{2\theta_i - 1} \leq 1$, as in Step 3 of the proof of Theorem 1. The argument is then repeated for the cones of \mathbb{R}^N corresponding to other orderings of the coordinates of ε . ■

If the interpretation of theorems 1 and 2 differ, their proofs are very similar.

We conclude this Section with a protocol implementing the liability profile just described. As in Proposition 2 this protocol is not the only one to do the job.

Proposition 3

Given N , a feasible liability profile θ as in Theorem 2 is implemented by the following protocol. For all $x, \delta \gg 0$:

$$\begin{aligned} w_i &= \delta_i \cdot \sum_{P(i, \sigma)} x_j - t_i \\ &= \{1 + (2\theta_i - 1)^2\} \sum_{j \in N \setminus i} \frac{\min\{\delta_j x_i, \delta_i x_j\}}{(2\theta_i - 1)^2 \delta_i x_i + (2\theta_j - 1)^2 \delta_j x_j} \cdot \delta_i x_i \end{aligned}$$

Proof

Fix i, j, x_i, δ_i , and two positive numbers α_i, α_j , then compute

$$\max_{x_j, \delta_j} \frac{\min\{\delta_j x_i, \delta_i x_j\}}{\alpha_i \delta_i x_i + \alpha_j \delta_j x_j} = \max_{x_j} \frac{x_i x_j}{\alpha_i x_i^2 + \alpha_j x_j^2} = \frac{1}{2\sqrt{\alpha_i \alpha_j}}$$

where the first equality is because the ratio increases in δ_j (resp. δ_i) if $\delta_j x_i \leq \delta_i x_j$ (resp. $\delta_j x_i \geq \delta_i x_j$). Applying this to the net cost $w_i(x, \delta)$ resulting from the protocol in the Proposition gives

$$\max_{x_{-i}, \delta_{-i}} w_i(x, \delta) = \left(1 + \frac{(2\theta_i - 1)}{2} \sum_{N \setminus i} \frac{1}{2\theta_j - 1}\right) \cdot \delta_i x_i = \theta_i \cdot \delta_i x_i$$

as claimed. ■

5 Queuing: off-line protocols

We return to the probabilistic server of Section 3. A *queuing problem* is a triple (N, x, τ) , where the set N of users is at most countable, and user i 's job of (positive) size x_i is released at time τ_i , $\tau_i \geq 0$. When N is infinite, we assume that the number of jobs released in any bounded interval $[0, a]$ is finite, so that the queue is finite at any point in time. In this Section and the next, the profile τ of release dates is entirely arbitrary, and our results are correspondingly fully general¹⁰.

Unlike in the scheduling problem, efficiency is now compatible with pre-emption. The only constraint is this: if job j preempts job i , namely the server starts processing job j while job i is either untouched or partially completed, then it must wait until job j is completed before returning to (or starting service on) job i . This is equivalent to selecting a certain priority ordering of N , and serving the highest priority among the jobs alive at any point in time (e.g., [9]).

As before the *slowdown* of job i in problem (N, x, τ) is $s_i(N, x, \tau) = \frac{y_i}{x_i}$, where y_i is the expected sojourn time of job i . Given N , a profile $\theta = (\theta_i, i \in N)$ is a *feasible liability profile* if there exists a protocol such that

$$s_i(N, x, \tau) \leq \theta_i \text{ for all } x \gg 0, \tau \geq 0 \text{ and all } i \in N \quad (7)$$

The key to this concept is the definition of a protocol, i.e., of the information that the server can use to prioritize the jobs. We shall look below at three very different informational assumptions, of which the first one is the subject of this brief Section. The server has full information at date 0 about the size and release dates of all jobs present and future. This is often called the *off-line* context, in which a protocol can be any mapping from a problem (N, x, τ) to a (deterministic or random) priority ordering of N .

In this context the feasible liability profiles θ are precisely the same as in the scheduling context, namely Proposition 1 and Theorem 1 are preserved word for word. Indeed fix any θ and assume there exists an off-line probabilistic (resp. deterministic) protocol guaranteeing property (7). The latter holds in particular for $\tau = 0$, thus $\sum_N \frac{1}{2\theta_i - 1} \leq 1$ follows by Theorem 1 (resp. $\sum_N \frac{1}{\theta_i} \leq 1$ by Proposition 1). Conversely, if $\sum_N \frac{1}{2\theta_i - 1} \leq 1$, we implement the liability θ by adapting the weighted quadratic (or serial) protocol of Proposition 2 as follows. The realizations of the variables Z_i determine the priority

¹⁰In Section 7 the release of new jobs follows a Poisson random process.

ordering of jobs and a job i is preempted by any job j that draws a higher priority. Clearly under this protocol for (N, x, τ) the sojourn time of any job is not larger than in $(N, x, 0)$ under the weighted quadratic (or serial) protocol of Proposition 2, hence the claim. In the case of deterministic protocols, we adapt similarly the *Earliest Deadline First* protocol of Proposition 1.

By the same argument, Theorem 2 extends to off-line protocols in queuing problems with cash transfers. In that context a queuing problem is a 4-uple (N, x, δ, τ) , where δ is the profile of waiting costs. To any such queuing problem we associate the scheduling problem $(N, x, \delta, 0)$, and write w_i^0 for the net waiting cost in this latter problem under a protocol such as the one in Proposition 3. Keep the efficient ordering σ used in $(N, x, \delta, 0)$ as the priority ordering dictating preemption in (N, x, δ, τ) . Note that σ is not necessarily efficient (may not minimize total waiting cost) in (N, x, δ, τ) ¹¹. Clearly the resulting net waiting cost w_i is not larger than w_i^0 , implying the claim.

6 Queuing: on-line protocols

We turn to the more interesting and more realistic *on-line* protocols, where the server has no information whatsoever about future jobs, so the protocol can only rely on the characteristics of live (released and not completed) jobs. In particular the server does not know which or how many jobs will be released in the future, hence the liability of job i cannot depend on its "name" within the unknown set N .

Check first that the on-line protocols commonly discussed in the literature do not offer a bounded liability. Consider first SRJF: if K jobs slightly shorter than job i are released almost immediately after τ_i , its slowdown reaches K , and this number cannot be bounded when the future is entirely opaque. A similar argument applies to LCFS and to PS, in which a burst of large jobs released just after τ_i increases the slowdown without bound. On the other hand, FCFS guarantees to any job a finite sojourn time independently of the future, yet the slowdown grows arbitrarily large with the size of the jobs in the existing queue. A similar argument applies to the queuing variant of RO, constructing the priority ordering on-line by giving to a new job an equal chance to each possible priority level among live jobs.

¹¹Writing \tilde{x}_i for the *remaining* service time of job i , the protocol serving at any time a job i such that $\frac{\tilde{x}_i}{\delta_i}$ is minimal among all live jobs, may still be inefficient, though less so than the protocol just described.

Yet it is feasible to achieve a finite liability for all jobs, irrespective of the size of live jobs and of the number, size and release dates of future jobs. This cap only depends upon the *number* of live jobs at release time, which we write as $r(i)$ for job i . Note that job i itself is counted in $r(i)$, and that $r(i)$ depends upon not only the problem (N, x, τ) but also the on-line protocol. In particular, $r(i)$ is random if the protocol is.

The protocol is inspired by the Fair Sojourn Processing ([8], [9]). At any point in time where m jobs are alive, the inefficient Processor Sharing protocol serves them all at the rate $\frac{1}{m}$. FSP runs the PS protocol *virtually*, and prioritizes the jobs according to their virtual completion date under PS (an earlier completion date means a higher priority). The result is an efficient deterministic protocol in which no sojourn times is longer than under PS.

In the variant of FSP establishing our next result, the server chooses once and for all a sequence $\theta_r, r = 1, 2, \dots$ of positive numbers such that $\sum_1^\infty \frac{1}{\theta_r} = 1$. At any point in time, the priority ordering of live jobs follows the ordering of their completion dates under the weighted version of PS where the processing rate of job i is proportional to $\theta_{r(i)}$. Thus one can think of weighted PS as a *virtual* protocol, used only to set priorities in the *real* (weighted) FSP protocol. We illustrate our protocol by an example with 4 jobs:

$$(x_1, \tau_1) = (10, 0); (x_2, \tau_2) = (6, 1); (x_3, \tau_3) = (4, 10); (x_4, \tau_4) = (1, 13)$$

To fix ideas we set $\theta_r = 2^r$, for all r . Starting at date 1, the virtual weighted PS processes jobs 1 and 2 at rates $2/3$ and $1/3$ respectively, because $\frac{\theta_1}{\theta_2} = 2$, thus, as long as these processing rates are maintained, job 1 is virtually completed at $\tau = 14.5$, before job 2 at $\tau = 19$. So in the real process job 1 has priority over job 2 and is completed at date 10.

At $\tau = 10$ job 3 is released, and $r(3) = r(2) = 2$. Note that we do not update $r(2)$ when job 1 exits. To determine the priority between jobs 2 and 3, check that under weighted PS, job 2 would have only 3 units left at date 10, and would be processed at the same rate as job 3 henceforth. Thus job 2 is virtually completed before job 3, hence it takes precedence over job 3 in the real process. The server processes only job 2 until $\tau = 13$. At that date three jobs are alive, and in the virtual weighted PS, only 1.5 units of job 2 are left (in reality, 3 units are left), because its rate of service was $1/3$ between $\tau = 1$ and $\tau = 10$, then $1/2$ between $\tau = 10$ and $\tau = 13$. Similarly only 2.5 units of job 3 are virtually left at date 13 (in fact, it is still intact).

The virtual rates of service starting at $\tau = 13$ are $2/5$, for jobs 1, 2 and $1/5$ for job 4. Thus job 2 has priority over job 4 and the latter over job 3. The real completion times are $\tau = 16$ for job 2, $\tau = 17$ for job 4, and $\tau = 21$ for job 3.

Theorem 3

Fix a non-decreasing sequence $\theta_r, r = 1, 2, \dots$ of positive numbers such that $\sum_1^\infty \frac{1}{\theta_r} = 1$. Then there exists an efficient on-line protocol implementing the liability $\theta_{r(i)}$ for job i :

$$s_i(N, x, \tau) \leq \theta_{r(i)} \text{ for all } x \gg 0, \tau \geq 0 \text{ and all } i \in N \quad (8)$$

Conversely such a protocol exists only if $\sum_1^\infty \frac{1}{\theta_r} \leq 1$.

Proof

Step 1. Suppose the liability $\theta_{r(i)}$ is implemented by some deterministic protocol (on-line or otherwise). Fix an integer n , a small positive number ε and consider the problem

$$N = \{1, 2, \dots, n\}; x_i = \frac{1}{\theta_i}, \tau_i = (i - 1)\varepsilon \text{ for all } i$$

Choose ε so that $n\varepsilon < \min_j x_j$. Thus all jobs are released before any job is completed, and $r(i) = i$ for all i , irrespective of the protocol. Suppose k is the job completed last, with sojourn time y_k . We have

$$y_k = \sum_N \frac{1}{\theta_i} - (k - 1)\varepsilon \leq y_k \leq \theta_k \cdot x_k = 1$$

from which the inequality $\sum_1^\infty \frac{1}{\theta_r} \leq 1$ follows in the limit.

Step 2. We fix the sequence θ_r such that $\sum_1^\infty \frac{1}{\theta_r} = 1$. To implement (8), we construct formally the protocol described before the Theorem. Let M be the set of live jobs at a given time, and consider a user i in M . Recall that $r(i)$ is the number of jobs that were alive *when job i was released*, and bears no relation to $m = |M|$. However we claim

$$\sum_M \frac{1}{\theta_{r(i)}} \leq 1 \quad (9)$$

To see this, label the jobs in M in the order of their release dates, say $\tau_1 \leq \tau_2 \leq \dots \leq \tau_m$. Note that jobs $1, \dots, i - 1$ are alive when i is released, therefore $r(i) \geq i$ for $i = 1, \dots, m$. So (9) follows by our assumptions on the sequence θ_r .

Let $\text{PS}[\theta]$ be the (inefficient) protocol dividing its time among the jobs in M in proportion to $\frac{1}{\theta_{r(i)}}$. Inequality (9) implies that the actual rate is no less than $\frac{1}{\theta_{r(i)}}$, therefore the sojourn time of i is at most $\theta_{r(i)} \cdot x_i$, and $\text{PS}[\theta]$ implements the liability (8) as announced.

Then $\text{FSP}[\theta]$ prioritizes the jobs according to their virtual completion times in $\text{PS}[\theta]$. This definition makes sense only if the relative ordering of any two jobs never changes when future jobs are released. Note that in $\text{PS}[\theta]$, the ratio of processing times for any two jobs i, j remains $\frac{\theta_i}{\theta_j}$ as long as they are both alive, so the ordering of their virtual completion times never changes either. Finally $\text{FSP}[\theta]$ implements the liability (8) because it is a Pareto improvement over $\text{PS}[\theta]$. ■

Step 1 in the proof of Theorem 3, shows its close relation with Proposition 1. The next question is whether or not an on-line probabilistic protocol can achieve a better liability function $r \rightarrow \theta_r$ than a deterministic one. I do not know the answer to this question. From Theorem 1 such a function has to meet $\sum_1^\infty \frac{1}{2\theta_r - 1} \leq 1$ (to see this use the same construction as in Step 1 of the above proof).

Remark 2

In the context with cash transfers and linear waiting costs, there is no hope to implement on-line a feasible liability of the type $r \rightarrow \theta_r$. Notice first that an on-line protocol cannot implement the fully efficient service ordering minimizing total waiting costs, because the computation of this ordering requires perfect forecast. The second best on-line protocol, from the point of view of efficiency, sets priorities at any time in the order of the ratios $\frac{\tilde{x}_i}{\delta_i}$, where \tilde{x}_i is the remaining service time of job i . Suppose we can define on-line transfers so as to implement the liability $r \rightarrow \theta_r$, and consider the following problem

$$\begin{aligned} (x_1, \delta_1, \tau_1) &= (a + \varepsilon, (1 - \varepsilon)a, 0), (x_2, \delta_2, \tau_2) = (1, 1, \varepsilon), \\ (x_i, \delta_i, \tau_i) &= (1, 1, i - 2 + \varepsilon) \text{ for } i = 3, \dots, n + 1 \end{aligned}$$

The server processes job 2 at $\tau = \varepsilon$, job i at $\tau = i - 2 + \varepsilon$, $i = 3, \dots, n + 1$, and finally job 1. Thus $r(i) = 2$ for $i = 2, \dots, n + 1$, and the total waiting cost

$$w_1 + \sum_2^{n+1} w_i = (1 - \varepsilon)a(a + \varepsilon + n) + n$$

should not exceed

$$\theta_1 \cdot \delta_1 x_1 + \sum_{i=2}^{n+1} \theta_2 \cdot \delta_i x_i = (1 - \varepsilon)a(a + \varepsilon)\theta_1 + n\theta_2$$

Letting ε go to zero we get

$$a(a + n) + n \leq a^2\theta_1 + n\theta_2 \text{ for all } a > 0, \text{ all } n = 1, 2, \dots$$

This is only possible if $\theta_2 = \infty$. A similar argument, omitted for brevity, shows $\theta_i = \infty$ for any $i \geq 3$ as well. This the trivial liability of FCFS.

7 On-line protocols with Poisson arrivals

We now assume that the release of new jobs follow a Poisson process with parameter λ , known to the server and the users. The size of all jobs, including future jobs remains unknown. The protocol is non anticipative, and only depends on λ and the characteristics of jobs currently alive. It is also *anonymous*, namely the only relevant characteristics of a job are its size and release date.

The arrival rate λ measures the degree of congestion of the queuing problem, and we seek a bound on the slowdown of a given job that depends *only* upon the size x of this job and λ . This is analogous to the anonymous liability $\frac{n+1}{2}$ in the scheduling problem (Theorem 1). The new feature, explained below, is that the liability depends also on own job size.

We consider a job of size x requested by user i , and released at a time where the arrival process has reached its steady state. In this Section we do not index the job size by the name of the user, because we only discuss anonymous protocols. Contrary to the previous Sections, x is now a positive number, not a vector in \mathbb{R}^N .

Given a protocol and an arbitrary sequence $\tilde{x} = (\dots, x_{-k}, \dots, x_{-1}, x_1, \dots, x_k, \dots)$ of job sizes for the predecessors (x_{-k}) and successors (x_k) of i 's job, we write $y(\lambda, x, \tilde{x})$ for the expected sojourn time of job i , and $s(\lambda, x, \tilde{x}) = \frac{y(\lambda, x, \tilde{x})}{x}$ for its expected slowdown. Note that the expectation only bears on the arrival process, in its steady state.

As we allow for jobs of arbitrary size, the queue may well be unstable (infinitely long in the steady state). This happens in particular if job sizes are drawn independently from a common distribution with mean no smaller than

$\frac{1}{\lambda}$. Therefore if all jobs have the same size $x, x \geq \frac{1}{\lambda}$, we cannot guarantee a bounded slowdown to all of them, and as a result the liability is infinite for this choice of λ and x .

In order to avoid topological difficulties arising when the sequence of job sizes contains subsequences converging to $\frac{1}{\lambda}$ from below, we will restrict attention to sequences \tilde{x} for which there exists a (small) positive number ε such that no job size x_k is in $[\frac{1}{\lambda} - \varepsilon, \frac{1}{\lambda}[$. We call *regular* a sequence with this property. The proof of Theorem 4 and Proposition 7 explains the role of this assumption.

We say that $\theta : \mathbb{R}_{++}^2 \rightarrow [1, \infty]$ is a *feasible liability function* if there exists an on-line anonymous protocol such that

$$s(\lambda, x, \tilde{x}) \leq \theta(\lambda, x) \text{ for all } \lambda, x > 0, \text{ and all regular sequences } \tilde{x} \quad (10)$$

We obtain first a simple lower bound on the feasible liability.

Proposition 4

If θ is a feasible liability function, we have

$$\theta(\lambda, x) \geq \frac{1}{1 - \lambda x}, \text{ if } x < \frac{1}{\lambda}, \theta(\lambda, x) = \infty \text{ if } x \geq \frac{1}{\lambda}$$

Proof.

It is enough to consider the case where all jobs are of identical size x . If $x \geq \frac{1}{\lambda}$, the queue is unstable, namely infinitely long in the steady state, therefore the sojourn time of some jobs is infinite, implying $\theta(\lambda, x) = \infty$. If $x < \frac{1}{\lambda}$, the mean sojourn time in any efficient protocol is $\frac{x}{1 - \lambda x}$ (e.g., [25]) and it can only be longer in an inefficient protocol, thus the first inequality follows. ■

The function $\theta^*(\lambda, x) = \frac{1}{1 - \lambda x}$, on $[0, \frac{1}{\lambda}[$, $= \infty$ on $[\frac{1}{\lambda}, \infty[$ is the *fair slowdown*, namely the slowdown a job can expect when all other jobs are of the same size. If I am not responsible for the size of the other jobs, this is a fair benchmark for the worst slowdown I could incur. The same idea has been used successfully in a variety of fair division problems (see [13], [14],[15]). Compute the utility I would receive when all other users have the same characteristics as me, and use it as a *lower bound* for my utility when they don't.

I conjecture that the idea applies to the M/G/1 queue as well, namely θ^* is a feasible liability function. In support of this conjecture, Theorem 4 offers a probabilistic protocol implementing the liability $(1.45)\theta^*$.

On the way to this result, we compute the liability functions of several familiar on-line protocols, and compare them to the fair slowdown.

Consider the protocols FCFS and LCFS. If we do not pay attention to the critical job size $\frac{1}{\lambda}$, the liability of any job is clearly infinite¹². Now fix a number $a, a < \frac{1}{\lambda}$, call "long" any job larger than a , "short" any other job, and consider the protocols FCFS[a], LCFS[a], in which all long jobs are systematically pushed back behind short jobs. So in FCFS[a] a new short job must only wait until all short jobs released before him are completed, and in LCFS[a] a short job is preempted only by short jobs released after him.

Proposition 5

Fix the arrival rate λ , and $a, 0 < a < \frac{1}{\lambda}$. The liability of LCFS[a] is

$$\theta^{LC}(\lambda, x) = \frac{1}{1 - \lambda a} \text{ if } x \leq a, \theta^{LC}(\lambda, x) = \infty \text{ if } x > a$$

The liability of FCFS[a] is

$$\theta^{FC}(\lambda, x) = (1 + \lambda a(\frac{a}{x} - 1)) \cdot \theta^{LC}(\lambda, x) \text{ if } x \leq a, \theta^{FC}(\lambda, x) = \infty \text{ if } x > a$$

In particular $\theta^{FC}(\lambda, x) \geq \theta^{LC}(\lambda, x)$, with equality only if $x \geq a$. Moreover $\lim_{x \rightarrow 0} \frac{\theta^{FC}(\lambda, x)}{\theta^{LC}(\lambda, x)} = +\infty$.

Proof

Under FCFS[a], the worst case for a job released at time t is that all jobs released earlier be of size a , in which case the expected length of the queue upon arrival of the new job is $\frac{\lambda a^2}{1 - \lambda a}$ (namely $\frac{a}{1 - \lambda a} - a$) hence

$$y = x + \frac{\lambda a^2}{1 - \lambda a} \implies s = 1 + \frac{\lambda a}{1 - \lambda a} \frac{a}{x} = (1 + \lambda a(\frac{a}{x} - 1)) \cdot \frac{1}{1 - \lambda a}$$

as claimed.

Consider next LCFS[a]. We write $W(q)$ for the worst expected time a queue made of short jobs and of total size q , takes to clear, where the maximization bears upon the size of future jobs, who will delay the completion of the

¹²If all previous jobs are longer than $\frac{1}{\lambda}$, the steady state queue is infinitely long and contains infinitely many jobs, therefore my sojourn under FCFS is unbounded. To prove the claim for LCFS: the probability that the next job released after mine shows up before my own job is completed is $1 - e^{-\lambda x} > 0$; if the size of that job grows arbitrarily large, so does my expected sojourn.

last job in the queue q . Clearly $W(q)$ obtains when all future jobs are of size a , and in this case $W(q) = \frac{q}{1-\lambda a}$. Indeed apart from the jobs in q , the server is idle (i.e., has no short job to process) on average during the fraction $(1 - \lambda a)$ of the time, and the initial job q will be processed only when the server is idle. ■

The truncated versions of LCFS and FCFS give their fair slowdown only to jobs of size exactly a . For job sizes between a and $\frac{1}{\lambda}$, the liability is infinite. To jobs no larger than a , LCFS[a] guarantees a multiple of the fair slowdown, but FCFS[a] does not:

$$\sup_{x \in [0, a]} \frac{\theta^{LC}(\lambda, x)}{\theta^*(\lambda, x)} = \frac{1}{1 - \lambda a}; \quad \lim_{x \rightarrow 0} \frac{\theta^{FC}(\lambda, x)}{\theta^*(\lambda, x)} = +\infty$$

Note the following tradeoff on a : if λa is small, θ^{LC} is close to θ^* for short jobs (e.g., for $\lambda a = 0.25$, $\theta^{LC} \leq (1.34)\theta^*$) but the set of such short jobs is small; if λa is close to 1, θ^{LC} becomes much larger than θ^* for very short jobs (e.g., for $\lambda a = 0.9$, $\theta^{LC} \leq 10\theta^*$).

We turn to the SRJF protocol, important because it minimizes total sojourn time.

Proposition 6

Fix the arrival rate λ . The liability of the SRJF protocol is bounded as follows

$$1 + \frac{\lambda x}{(1 - \lambda x)^2} \leq \theta^{SR}(\lambda, x) \leq \frac{1}{(1 - \lambda x)^2} \text{ if } x < \frac{1}{\lambda}, \quad \theta^{SR}(\lambda, x) = \infty \text{ if } x \geq \frac{1}{\lambda}$$

The longer and related proofs of Proposition 6 and Theorem 4 are in the Appendix.

Note that the bounds for θ^{SR} are fairly tight: they coincide for $x = 0$, their ratio never exceeds $\frac{4}{3}$ and goes to 1 as x approaches $\frac{1}{\lambda}$.

Comparing the liabilities offered by SRJF and LCFS[a], we see first that their worst performance on $[0, a]$, relative to the fair slowdown, are essentially identical:

$$\sup_{x \in [0, a]} \frac{\theta^{SR}(\lambda, x)}{\theta^*(\lambda, x)} \simeq \frac{1}{1 - \lambda a} \text{ for } a \rightarrow \frac{1}{\lambda}$$

However the supremum is reached for $x = 0$ in the case of LCFS[a], and for $x = a$ in the case of SRJF. The latter behaves much better than the former for short jobs, as $\lim_{x \rightarrow 0} \frac{\theta^{SR}(\lambda, x)}{\theta^*(\lambda, x)} = 1$; the reverse is true for job sizes near a .

In the protocols discussed so far, the liability function does not approximate the fair slowdown over the entire interval $[0, \frac{1}{\lambda}]$. Our next family of *probabilistic priority* protocols achieves just that. Like the parametric protocols of Section 3, each such protocol is defined by a family of cumulative distribution functions $F_x(q)$ on $[0, +\infty[$, one for each job size $x, 0 < x < \frac{1}{\lambda}$.

Given such a family F of c.d.f.s, we denote by $P(F)$ the following protocol. Any job $x, x \geq \frac{1}{\lambda}$, is pushed at the back of the queue upon its release. For each new job $x, x < \frac{1}{\lambda}$, draw a random variable Z according to F_x , and use its realization to insert job x in the queue. More precisely suppose upon the release of x the queue is $\{x_1, x_2, \dots, x_K\}$, where x_k is the remaining size of the job ranked k . The new queue is

$$\{x_1, \dots, x_k, x, x_{k+1}, \dots, x_K\} \text{ if } \sum_1^k x_j \leq Z < \sum_1^{k+1} x_j, \text{ for } k = 0, 1, \dots, K$$

with the convention $x_{K+1} = \infty$. In particular the new job is first in the queue if $Z < x_1$, and last if $\sum_1^K x_j \leq Z$. Note also that a "short" job ($x < \frac{1}{\lambda}$) does not care about future "long" jobs, but long jobs released earlier matter because they are part of the queue in which it is inserted.

Theorem 4

It is possible to choose the c.d.f.s $\{F_x, 0 < x < \frac{1}{\lambda}\}$ in such a way that the liability θ of the probabilistic priority protocol satisfies

$$\theta(\lambda, x) \leq (1.45) \cdot \theta^*(\lambda, x) \text{ for all } x$$

One such choice is

$$F_x(q) = \min\left\{\frac{1}{x} \cdot \frac{2q}{2\lambda \cdot q + \lambda \cdot x + 1.25}, 1\right\} \text{ for all } q \geq 0,$$

implementing the liability function

$$\theta(\lambda, x) = f(\lambda x), \text{ where } f(u) = \left(\frac{1}{4} + \frac{u}{5}\right) \left(\frac{3.25 - u}{1 - u} - \frac{\log(1 - u)}{2u}\right)$$

Figure 1 depicts the ratio $\frac{\theta}{\theta^*}$ as a function of $u = \lambda x$, for the above protocol.

In the proof of Theorem 4, we describe a rich family of probabilistic priority protocols for which the liability function can be estimated numerically:

Proposition 7 below. It is an open question to find out the optimal choice of the c.d.f.s $\{F_x, 0 < x < \frac{1}{\lambda}\}$, namely that for which the upper bound on the ratio $\frac{\theta}{\theta^*}$ is the smallest. Perhaps some choice of F implements the fair slowdown $\theta = \theta^*$.

Another open question, of a more technical nature, is the role of our regularity assumption on sequences \tilde{x} . Does Theorem 4 hold if we request inequality (10) to hold for all sequences \tilde{x} ?

References

- [1] Bansal, N. and Harchol-Balter, M. (2001). "Analysis of SRPT Scheduling: Investigating Unfairness," in *Proc. ACM Sigmetrics '01*.
- [2] Bansal, N. and Wierman, A. (2002). "Competitive Analysis of M/GI/1 Queueing Policies," mimeo, Carnegie-Mellon University.
- [3] Bender, M., Chakrabarti, S., and Muthukrishnan, S. (1998). "Flow and Stretch Metrics for Scheduling Continuous Job Streams," in *Proc. 9th ACM-SIAM Symposium on Discrete Algorithms*.
- [4] Chun, Y. (2004). "Consistency and Monotonicity in Sequencing Problems, mimeo, Seoul National University.
- [5] Demers A., S. Keshav and S. Shenker (1990) "Analysis and simulation of a fair queuing algorithm", *Internetworking: Research and Experience*, 1, 3-26.
- [6] Demko and T. Hill (1988). "Equitable distribution of indivisible objects", *Mathematical Social sciences*, 16,2, 145-58.
- [7] Dubins, L.F. (1977). "Group Decision Devices," *Amer. Math Monthly*, May, 350-356.
- [8] Friedman, E.J. and Henderson, S.G. (2003). "Fairness and Efficiency in Web Server Protocols," in *Proceedings of the 2003 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, 229-237, ACM Press.
- [9] Friedman, E.J., Henderson, S.G., and G. Hurley, (2004) Minimizing Mean Response Time Subject to Fairness, *mimeo*, Cornell University

- [10] Harchol-Balter, M., Bansal, N., Schroeder, B., and Agrawal, M. (2001). "Size-based Scheduling to Improve Web Performance," mimeo, Carnegie-Mellon University.
- [11] Maniquet, F. (2003). "A Characterization of the Shapley Value in Queueing Problems," *Journal of Economic Theory*, **109**, 1, 90-103.
- [12] Mitra, M. (2001). "Mechanism Design in Queueing Problems," *Economic Theory*, **17**, 277-305.
- [13] Moulin, H. (1990). "Uniform Externalities: Two Axioms for Fair Allocation," *Journal of Public Economics*, 305–326.
- [14] Moulin, H. (1991). "Welfare Bounds in the Fair Division Problem," *Journal of Economic Theory*, **54**, 2, 321–337.
- [15] Moulin, H. (1992). "Welfare Bounds in the Cooperative Production Problem," *Games and Economic Behavior*, **4**, 373–401.
- [16] Moulin, H. (2004). "Split-proof probabilistic scheduling," mimeo, Rice University.
- [17] Queyranne, M. (1993). "Structure of a Simple Scheduling Polyhedron," *Mathematical Programming*, **58**, 263-285.
- [18] Shapley, L. (1971). Core of Convex Games, *International Journal of game Theory*, **1**, 11-26.
- [19] Smith, W., (1956). "Various Optimizers for Single-Stage Production", *Naval Res. Logistics Quarterly* **3**, 59-66
- [20] Steinhaus, H. (1948). "The Problem of Fair Division," *Econometrica* **16**, 101-104.
- [21] Suijs, J. (1996). "On Incentive Compatibility and Budget Balancedness in Public Decision Making," *Economic Design*, **2**, 193-209.
- [22] Thomson, W., and Varian, H. (1985). "Theories of Justice Based on Symmetry," in *Social Goals and Social Organizations* (Hurwicz *et al.*, Eds.). Cambridge: Cambridge University Press

- [23] Wierman, A. and Harchol-Balter, M. (2003). "Bounds on a Fair Policy with Near Optimal Performance," submitted.
- [24] Wierman, A. and Harchol-Balter, M. (2003). "Classifying Scheduling Policies with Respect to Unfairness in an M/GI/1," in *Proceedings of ACM Sigmetrics Conference on Measurement and Modeling of Computer Systems*.
- [25] Wolff, R. W. (1989). *Stochastic modeling and the theory of queues*, Prentice-Hall series in Industrial and Systems Engineering, New Jersey

8 Appendix: Proposition 7; proofs of Proposition 6, Theorem 4

8.1 Proof of Proposition 6

Step 1. We compute first the liability of the protocol *Shortest Job First* (SJF), namely the modification of SRJF where at any time the server processes (one of) the job that was shortest *at release time*.

Fix a job of size x , $x < \frac{1}{\lambda}$, at release time, and consider a queue of total length q in which job x is last ($q < x$ is possible if job x is partially completed and is alone in the queue) and all jobs before x in the queue were shorter than x at their release time. For such a queue, define $W(q|x)$ as the worst expected time it takes the queue to clear, where the maximization bears upon the size of future jobs (as in the proof of Proposition 5). This is well defined because subsequent jobs preempt job x if and only if they are shorter than x , so we do not need to specify the initial length of jobs ahead of job x . Clearly $W(q|x)$ is achieved when all future jobs are of size barely below x , therefore W satisfies the following functional equation, where for simplicity we omit x in $W(q|x)$:

$$W(q) = \exp(-\lambda q)q + \int_0^q \lambda \exp(-\lambda t) \{t + W(q - t + x)\} dt$$

$$\Leftrightarrow W(q) = \frac{1 - \exp(-\lambda q)}{\lambda} + \lambda \exp(-\lambda q) \int_0^q \exp(\lambda t) W(t + x) dt \text{ for all } q \quad (11)$$

Routine computation shows that the function

$$W(q) = \frac{q}{1 - \lambda x} \quad (12)$$

is a solution of (11). To see that it is the only one, observe that W grows no more than linearly, because when all future jobs are of size x , the server is idle a positive fraction of the time. Therefore for any positive number μ , the integral

$$\|W\| = \int_0^\infty |W(q)| \exp(-\mu q) dq$$

is a well defined norm for W . In the space of continuous functions for which the above integral converges, consider the linear mapping Φ

$$\Phi(W)(q) = \lambda \exp(-\lambda q) \int_0^q \exp(\lambda t) W(t + x) dt$$

(note that Φ depends upon the job size x , that remains fixed throughout). We claim that, for an appropriate choice of μ , Φ is contracting for the corresponding norm $\|\Phi\| = \sup \frac{\|\Phi(W)\|}{\|W\|}$:

$$\begin{aligned} \|\Phi(W)\| &= \int_0^\infty \lambda \exp(-\lambda q) \left\{ \left| \int_0^q \exp(\lambda t) W(t + x) dt \right| \right\} \exp(-\mu q) dq \\ &\leq \int \int_{0 \leq t \leq q < \infty} \lambda \exp(-(\lambda + \mu)q) \exp(\lambda t) |W(t + x)| dq dt \\ &= \frac{\lambda}{\lambda + \mu} \int_0^\infty |W(t + x)| \exp(-\mu t) dt = \frac{\lambda}{\lambda + \mu} \exp(\mu x) \int_x^\infty |W(t)| \exp(-\mu t) dt \\ &\leq \frac{\lambda}{\lambda + \mu} \exp(\mu x) \|W\| \end{aligned}$$

Now for μ small enough, $\lambda x < 1$ implies $\frac{\lambda}{\lambda + \mu} \exp(\mu x) < 1$, establishing that Φ is contracting. As the above functional space is complete for this norm, we conclude that the equation $W = U_0 + \Phi(W)$ has at most one solution for any constant U_0 .

Having determined W ((12)), we now compute the liability of job x under SJF. Upon its release, it will be pushed behind all jobs still alive and originally shorter than x . Thus the longest queue ahead of job x when it is released,

obtains when *all* past jobs were barely shorter than x . In that case the expected length of this queue (including job x) is $q = \frac{x}{1-\lambda x}$. Thus the worst expected delay for job x is $W(\frac{x}{1-\lambda x}) = \frac{x}{(1-\lambda x)^2}$, and the liability of SJF is $\frac{1}{(1-\lambda x)^2}$.

Step 2. In this step we show that the liability of SRJF is not worse than that of SJF. In order to maximize the steady-state sojourn time of a given job of size x , $x < \frac{1}{\lambda}$, it is enough to look at the case where all past and future jobs are no larger than x . This is clear for future jobs, and for past jobs the only qualification to this statement comes if there is a single other job, of size x_{-1} , alive when job x is released. Then the worst case is when what is left of x_{-1} at the release of x is just below x . But this only occurs when there is a single other job alive, and the adversary "Nature" can predict the exact release date of job x , which is ruled out by our definition of the liability function ((10)).

Now fix a sequence $\tilde{x} = (\dots, x_{-k}, \dots, x_{-1}, x_1, \dots, x_k, \dots)$ of sizes all below x for the other jobs, and arbitrary release dates. Compare then the sojourn time of job x under SRJF and under SJF: under the latter, all other jobs take precedence over x , but this is not necessarily the case under the former; as soon as job x is partially processed, it concedes priority to fewer jobs under SRJF. In view of Step 1, this proves the upper bound on $\theta^{SR}(\lambda, x)$.

Step 3. Finally we prove the lower bound on $\theta^{SR}(\lambda, x)$ by computing the expected slowdown when all past and future jobs are barely shorter than x . Note that this is not the optimal choice of the sequence \tilde{x} in (10): future jobs expect to be compared to a partially completed job "x", hence the optimal choice involves somewhat shorter future jobs. Once again we compute the function W defined as in Step 1, solving now a slightly different functional equation, because as soon as the server starts processing job x it is not interrupted by any subsequent job:

$$\begin{aligned} W(q) &= \exp(-\lambda(q-x))q + \int_0^{q-x} \lambda \exp(-\lambda t) \{t + W(q-t+x)\} dt \text{ if } q \geq x \\ &= q \text{ if } 0 \leq q \leq x \end{aligned}$$

Changing the unknown function to $\tilde{W}(\tilde{q}) = W(\tilde{q} + x) - x$, we find that \tilde{W} satisfies the same equation (11), therefore

$$W(q) = x + \frac{q-x}{1-\lambda x} \text{ if } q \geq x$$

Now, as in Step 1, the queue ahead of (and including) job x when it is released has expected length $q = \frac{x}{1-\lambda x}$, so the expected slowdown is

$$s = \frac{W\left(\frac{x}{1-\lambda x}\right)}{x} = 1 + \frac{\lambda x}{(1-\lambda x)^2}$$

This concludes the proof.

8.2 Proposition 7

The arrival rate λ of new jobs is fixed throughout.

In the statement of Proposition 7, we are given a real valued function W , differentiable on $[0, +\infty[$ with the following properties:

$$W(0) = 0; W'(0) = 1; \text{ and for all } x, 0 < x < \frac{1}{\lambda},$$

$$R(q, x) = \frac{W'(q) - 1}{W(q+x) - W(q)} \text{ increases strictly in } q; \lim_{q \rightarrow \infty} R(q, x) = \frac{1}{x} \quad (13)$$

Given such a function W , the equation in q , $R(q, x) = \lambda$ has a unique solution $q(x)$ for all $x, 0 < x < \frac{1}{\lambda}$, and we further assume

$$\lim_{x \rightarrow \frac{1}{\lambda}} q(x) = \infty \quad (14)$$

Proposition 7

For any function W satisfying (13) and (14), consider the probabilistic priority protocol $P(F)$ where

$$F_x(q) = \min\left\{\frac{R(q, x)}{\lambda}, 1\right\} \text{ for } x > 0, q \geq 0 \quad (15)$$

Then $W(q)$ is the worst expected time to clear a queue of size q under $P(F)$, and this protocol implements the liability function

$$\begin{aligned} \theta(\lambda, x) &= \int_0^{q(x)} W(q+x) dF_x(q) \\ &= \frac{1}{x} \left\{ W(x) + \int_0^{q(x)} \left(1 - \frac{R(q, x)}{\lambda}\right) W'(q+x) dq \right\} \end{aligned} \quad (16)$$

Proof

Our assumptions on W ensure that equation (15) defines for all λ and all x a continuous c.d.f. F_x with support $[0, q(x)]$.

We introduce first some notations . Let $X = \mathbb{R}_{++}^{\mathbb{Z} \setminus 0}$ be the set of sequences \tilde{x} where $\mathbb{R}_{++} =]0, \infty[$. Let X^r be the subset of regular sequences, namely there exists $\varepsilon, \varepsilon > 0$, such that $x_k \notin [\frac{1}{\lambda} - \varepsilon, \frac{1}{\lambda}[$ for all k . For all $a \in]0, \frac{1}{\lambda}[$, let $X(a)$ be the subset of X^r such that $x_k \notin [a, \frac{1}{\lambda}[$ for all k . The sets $X(a)$ are nested and their union is X^r . Let $P(a, F)$ be the " a -truncation" of $P(F)$ where all jobs longer than a are systematically pushed at the end of the queue. Thus in $P(a, F)$ we treat a job $x, x > a$, exactly as we treat jobs $x, x \geq \frac{1}{\lambda}$ in $P(F)$. *Step 1.* We write $V(q)$ (resp. $V^a(q)$) for the worst expected time to clear a queue of length q under $P(F)$ (resp. $P(a, F)$), where the maximization bears, as usual, on the size of future jobs. In Steps 2 and 3 below we show that $V^a(q)$ is the following function

$$\begin{aligned} V^a(q) &= W(q) \text{ if } 0 \leq q \leq q(a) \\ &= W(q(a)) + \frac{q - q(a)}{1 - \lambda a} \text{ if } q(a) \leq q \end{aligned} \quad (17)$$

Here we check that (17) implies Proposition 7. Write $V(q; \tilde{x})$ (resp. $V^a(q; \tilde{x})$) for the expected time under $P(F)$ (resp. under $P(a, F)$) to clear a queue of size q given a sequence \tilde{x} of (future) jobs sizes. Clearly $V^a(q; \tilde{x}) = V(q; \tilde{x})$ whenever $\tilde{x} \in X(a)$, therefore

$$V(q) = \sup_{X^r} V(q; \tilde{x}) = \sup_{a \in]0, \frac{1}{\lambda}[} \sup_{X(a)} V(q; \tilde{x}) = \sup_{a \in]0, \frac{1}{\lambda}[} V^a(q) = W(q)$$

where the right hand equality follows from (17) and assumption (14).

We show next that $P(F)$ implements the liability function (16). Fix any $x \in [0, \frac{1}{\lambda}[$, and choose a^* such that $q(a^*) \geq x + q(x)$, which is possible by (14). In particular $x \leq a^*$. We claim now

$$\sup_{X(a)} y(\lambda, x, \tilde{x} | P(a, F)) \leq \theta(\lambda, x)x \text{ for all } a, a^* \leq a < \frac{1}{\lambda}, \quad (18)$$

(recall that $y(\lambda, x, \tilde{x})$ is the expected sojourn of job x given \tilde{x}). Assume for a moment the claim holds. By construction $x \leq a^* \leq a$, therefore $y(\lambda, x, \tilde{x})$ is the same under $P(a, F)$ and under $P(F)$ and (18) holds as well when we

replace $P(a, F)$ by $P(F)$. Moreover for any $\tilde{x} \in X^r$ we can choose a such that $a^* \leq a < \frac{1}{\lambda}$ and $\tilde{x} \in X(a)$, completing the proof that $P(F)$ implements θ .

It remains to prove property (18). Fix a as announced there, in particular $x \leq a$ and $x + q(x) \leq q(a)$. Upon its release, job x draws Z with distribution F_x and is placed at the back of a queue of length at most Z , resulting in a worst expected sojourn no larger than $V^a(q + Z)$. In view of (17) this implies

$$\sup_{X(a)} y(\lambda, x, \tilde{x} | P(a, F)) \leq \int_0^{q(x)} V^a(q + x) dF_x(q) = \int_0^{q(x)} W(q + x) dF_x(q)$$

Finally we use (15) and an integration by parts to compute

$$\begin{aligned} \int_0^{q(x)} W(q + x) dF_x(q) &= W(q + x) - \int_0^{q(x)} F_x(q) W'(q + x) dq \\ &= W(x) + \int_0^{q(x)} \left(1 - \frac{R(q, x)}{\lambda}\right) W'(q + x) dq \end{aligned}$$

Step 2. We call W^a the function on the right hand side of (17), and in this step we show $V^a \geq W^a$ by proving that, for all q , the expected time $U^a(q)$ to clear a queue of size q when *all* future jobs are of size a is $U^a(q) = W^a(q)$.

If a new job shows up when the queue size is q , the clearing time increases to $U^a(q + a)$ with probability $F_a(q)$ and is otherwise unaffected, therefore U^a satisfies the following equation

$$U^a(q) = \exp(-\lambda q)q + \int_0^q \lambda \exp(-\lambda t) \{t + H(U^a)(q - t)\} dt \Leftrightarrow$$

$$U^a(q) = \frac{1 - \exp(-\lambda q)}{\lambda} + \lambda \exp(-\lambda q) \int_0^q \exp(\lambda t) H(U^a)(t) dt, \text{ all } q \geq 0 \quad (19)$$

where H is the following (linear) operator for arbitrary U :

$$H(U)(q) = F_a(q)U(q + a) + (1 - F_a(q))U(q), \text{ all } q \geq 0$$

We solve this equation in two steps.

Case 1: $q \geq q(a)$

In this case $H(U)(q) = U(q + a)$, so that equation (19) is identical to (11), except for the restriction on the domain of q . We claim that its unique solution is

$$U^a(q) = U^a(q(a)) + \frac{q - q(a)}{1 - \lambda a} \text{ for all } q \geq q(a) \quad (20)$$

To prove the claim, we note that equation (19) implies

$$\int_0^{q(a)} \exp(\lambda t) U^a(t + a) dt = \frac{\exp(\lambda q(a))}{\lambda} \left\{ U^a(q(a)) - \frac{1 - \exp(-\lambda q(a))}{\lambda} \right\}$$

Then for $q \geq q(a)$, we replace the integral in the RHS of (19) by

$$\int_0^q \exp(\lambda t) U^a(t + a) dt = \int_0^{q(a)} \exp(\lambda t) U^a(t + a) dt + \int_{q(a)}^q \exp(\lambda t) U^a(t + a) dt$$

Combining the two above equations with (19), a straightforward computation gives:

$$\begin{aligned} U^a(q) - U^a(q(a)) &= \frac{1 - \exp\{-\lambda(q - q(a))\}}{\lambda} \\ &+ \lambda \exp\{-\lambda(q - q(a))\} \int_{q(a)}^q \exp\{\lambda(t - q(a))\} (U^a(t + a) - U^a(q(a))) dt \end{aligned}$$

which is precisely the same equation as (11) up to the change of unknown function $U^*(q) = U^a(q + q(a)) - U^a(q(a))$. The same argument as in Step 1 of the proof of Proposition 6 shows that U^a given by (20) is the unique solution of (19) on $[q(a), \infty[$.

Case 2: $0 \leq q \leq q(a)$

In equation (19) the integral involves values taken by U^a at $q + a$, yet by Case 1 we can regard the unknown in equation (19) as a continuous function defined on $[0, q(a)]$, with the convention that it extends continuously beyond this interval as the straight line with slope $\frac{1}{1 - \lambda a}$. Now we check that the operator Ψ , defined as follows for continuous functions U on $[0, q(a)]$:

$$\Psi(U)(q) = \lambda \exp(-\lambda q) \int_0^q \exp(\lambda t) H(U)(t) dt$$

is contracting for the supremum norm $\|U\|_\infty = \sup_{[0, q(a)]} |U(q)|$. Compute first the norm of H . For all U^1, U^2

$$\begin{aligned} |H(U^1)(q) - H(U^2)(q)| &\leq \sup_{[0, q(a)]} \{ |(U^1 - U^2)(q+a)|, |(U^1 - U^2)(q)| \} \\ &= \sup_{[0, q(a)]} |(U^1 - U^2)(q)| = \|U^1 - U^2\|_\infty \end{aligned}$$

where the equality follows from the canonical extension of both functions beyond $q(a)$. Then

$$\begin{aligned} \Psi(U^1 - U^2)(q) &\leq \{ \lambda \exp(-\lambda q) \int_0^q \exp(\lambda t) dt \} \|U^1 - U^2\|_\infty \\ &= (1 - \exp(-\lambda q)) \|U^1 - U^2\|_\infty \leq (1 - \exp(-\lambda q(a))) \|U^1 - U^2\|_\infty \end{aligned}$$

concluding the proof that Ψ is contracting, and that equation (19) has a unique solution on $[0, q(a)]$. It remains to check that W is this solution. Compute for all $q \in [0, q(a)]$

$$H(W)(q) = W(q) + F_a(q)(W(q+a) - W(q)) = W(q) + \frac{W'(q) - 1}{\lambda}$$

A routine computation in (19) establishes the claim.

Step 3. We show $V^a \leq W^a$. We compute the worst expected time $T^a(q)$ to clear a queue of size q when the adversary "Nature" chooses the size of each future job as a function of the remaining queue size. This is more than it is allowed to do under our definition of $y(\lambda, x, \tilde{x})$, therefore $V^a \leq T^a$. We prove $T^a = W^a$.

The functional equation satisfied by T^a is

$$T^a(q) = \exp(-\lambda q)q + \int_0^q \lambda \exp(-\lambda t) \{t + K(T^a)(q-t)\} dt \Leftrightarrow$$

$$T^a(q) = \frac{1 - \exp(-\lambda q)}{\lambda} + \lambda \exp(-\lambda q) \int_0^q \exp(\lambda t) K(T^a)(t) dt, \text{ all } q \geq 0 \quad (21)$$

where K is the following (non linear) operator for arbitrary T :

$$K(T)(q) = \sup_{x \in [0, a]} \{F_x(q)T(q+x) + (1 - F_x(q))T(q)\} \text{ for all } q \geq 0$$

The proof parallels that of Step 2.

We consider first (21) for $q \geq q(a)$. Then $K(T)(q) = T(q + a)$ provided T is non decreasing, which is clearly true for T^a , so we are back to Case 1 in the previous Step. We conclude that T^a satisfies (20).

Next we consider (21) on $[0, q(a)]$, using the same convention as in Case 2 above: we view the unknown as a function on $[0, q(a)]$ canonically extended beyond this interval by (20). We define an operator Θ

$$\Theta(T)(q) = \lambda \exp(-\lambda q) \int_0^q \exp(\lambda t) K(T)(t) dt$$

and show as in case 2 that it is contracting for the supremum norm $\|T\|_\infty$. We have

$$|K(T^1)(q) - K(T^2)(q)| \leq \sup_{[0, q(a)+a]} |(T^1 - T^2)(q)| = \sup_{[0, q(a)]} |(T^1 - T^2)(q)|$$

where the last equality follows the extension property. The proof that Θ is contracting, i.e., $\|\Theta\| \leq 1 - \exp(-\lambda q(a))$, follows as above.

Finally we check that the function W satisfies (21) on $[0, q(a)]$. Fix q in this interval; the definition of F ((15)) implies

$$\begin{aligned} F_x(q)(W(q+x) - W(q)) &= \frac{W'(q) - 1}{\lambda} \text{ if } q(x) \geq q \\ &\leq \frac{W'(q) - 1}{\lambda} \text{ if } q(x) \leq q \\ \implies K(W)(q) &= W(q) + \frac{W'(q) - 1}{\lambda} \text{ for all } q \in [0, q(a)] \end{aligned}$$

and the desired conclusion follows exactly as in Case 2.

8.3 Proof of Theorem 4

Fix a number $b, b > 0$, and define $W(q) = q + b\lambda q^2$. We let the reader check that W meets properties (13) and (14):

$$\begin{aligned} \frac{W'(q) - 1}{W(q+x) - W(q)} &= \frac{2b\lambda q}{x(1 + b\lambda x + 2b\lambda q)} \\ q(x) &= \frac{\lambda x}{1 - \lambda x} \cdot \frac{1 + b\lambda x}{2b\lambda} \end{aligned}$$

Thus Proposition 7 applies and the protocol $P(F)$ with

$$F_x(q) = \min\left\{\frac{2bq}{x(1 + b\lambda x + 2b\lambda q)}, 1\right\} \quad (22)$$

implements the liability function

$$\begin{aligned} \theta(\lambda, x) &= \int_0^{q(x)} W(q+x) dF_x(q) = \frac{2b\lambda(1+b\lambda x)}{\lambda x} \int_0^{q(x)} \frac{(x+q)(1+b\lambda x+b\lambda q)}{(1+b\lambda x+2b\lambda q)^2} dq \\ &= (1+b\lambda x) \cdot \left\{ \frac{1+2b-b\lambda x}{4b(1-\lambda x)} + \frac{|\log(1-\lambda x)|}{2\lambda x} \right\} \end{aligned}$$

Setting $\lambda x = u$, the ratio $\frac{\theta}{\theta^*}$ of the above liability to the fair slowdown is:

$$\frac{\theta}{\theta^*}(u) = \frac{1}{2}(1+bu)\left(1 + \frac{1}{2b} - \frac{u}{2} + \frac{|\log(1-u)|(1-u)}{u}\right) \text{ for } 0 \leq u \leq 1$$

One checks easily that this function is concave on $[0, 1]$, increasing from $1 + \frac{1}{4b}$ at $u = 0$, decreasing to $\frac{(b+1)^2}{4b}$ at $u = 1$. To minimize the maximal value of the ratio $\frac{\theta}{\theta^*}$ on $[0, 1]$, the optimal choice of b is near $b = 0.8$. Figure 1 depicts the ratio in this case, with its maximal value of 1.449. The corresponding family of c.d.f.s (22) is precisely the one in the statement of Theorem 4.

The rationale for choosing a quadratic function W is clear: it yields simple formulas for $F_x, q(x)$ and $\theta(\lambda, x)$. Happily, it also delivers a fairly good approximation of the fair slowdown. But there is no reason to suspect that another choice of W would not improve upon this approximation, even perhaps give an exact implementation of the fair slowdown. The difficulty is that for most other choices of W , we cannot get a closed form for F_x or $q(x)$, and the optimization of W is a serious numerical challenge.