Mechanism Design for a Service Provided in the Cloud

Richard Weber, University of Cambridge

Talk to Management Science Group at LSE

5.0pm, 27 February, 2013

Abstract

- It is proposed to provide a new web service that allows its users to share digital content (such as photos, files, etc) in the cloud. We ask potential users to declare their private values for the service, and on that basis we decide whether or not to provide it. If we do, some persons are told the login password, others are excluded.
- Ideally, no person wishes to change her declaration ex-post (just as occurs under the well-known Vickrey auction in its context). Unfortunately, this goal is too ambitious. A canonical optimal mechanism design relaxes some ex-post constraints, but it is complicated, and impractical in that persons who are excluded from accessing service may be required to provide some of its content — a rule which may be unrealistic to enforce!
- We will describe how to construct a mechanism that achieves welfare as great as the canonical design, has the desired ex-post budget balance, and respects the requirement that no excluded person should have to provide content.
- There is a simple design which is asymptotically optimal as the number of agents becomes large.

1 Mechanism design for a photo sharing library

Peer i shares p_i photos.
The benefit to peer i is $\theta_i u(p_1, \dots, p_n)$.
In this talk $u = 1$ (or 0) as $p_1 + \cdots + p_n \geq c$ (or $< c$).
Think of p_i as a payment made in kind. The sum of the payments must cover a cost c.

2 Mechanism design for an occasional service

2.1 Definitions

It costs c to provide a service. Some days it is provided and some days not.

Service might be ‘access to a web library of photos’.

or 'new research papers read and summarized for others to read'

Cost to each agent is her total dropbox space, occupied by photos, and the time she must take to upload them,

or the numbers of papers she has to read and summarize.

2.2 Private values

Each day we ask each of n agents to state some private information about how much they would value access to the service that day.

- Suppose private values are $\theta_1, \dots, \theta_n$.

  Assume (for simplicity of exposition) ex-ante values are i.i.d. $U[0,1]$.

- Agent i states $\theta_i = t_i$ and agrees to pay a fee. Her expected fee is $p(t_i)$.

  Perhaps she promises to upload $p(t_i)$ photos.

- On the basis of the stated $t_1, \dots, t_n$, we then decide either

  - not to provide the service (say if the t_i are smallish), or
  - provide the service, and selectively give access to some subset of the agents (those who have stated large values of t_i).

3 Requirements

We wish to maximize social welfare subject to 3 requirements:

(i) It should optimal for an agent truthfully to state $t_i = \theta_i$.

Revelation principle says this is not a restrictive requirement.

We will learn $\theta = (\theta_1, \dots, \theta_n)$.

We say that the mechanism design should be **ex-ante individually rational** and **ex-ante incentive compatible**.

(ii) Service is provided by payments in kind so the sum of the collected fees must exactly equal c (if the service provided) and no fees should be collected if not provided. This is called **ex-post budget balance**.

(iii) It would be unfair (and hard to enforce) to ask an agent to provide dropbox space or photos for the benefit of others if she herself is not allowed to access the photos!

So in cases that the service is provided, but some agents are excluded, the excluded agents should not be charged any fee (or compensated) (**no ex-post regrets for excluded agents**).

$$1_i(\theta) = 0 \implies p_i(\theta) = 0.$$

4 Notation

Some notation to describe the mechanism:

$1(\theta) = 1$ (or 0) as service is (or is not) provided,

$1_i(\theta) = 1$ (or 0) as access to service is (or is not) provided to agent i,

Clearly need $1_i(\theta) \leq 1(\theta)$.

$\pi = E[1(\theta)]$, (probability service provided)

$\pi_i(\theta_i) = E[1_i(\theta) \mid \theta_i]$, (probability service provided to agent i when she declares θ_i)

$p_i(\theta) =$ fee charged to agent i when agents declare $\theta = (\theta_1, \dots, \theta_n)$.

5 Standard mechanism design theory

Standard MD theory tells us what to do if we can substantially relax both (ii) and (iii).

Suppose (ii) is relaxed so we only require **ex-ante** budget balance:

$$\text{expected cost} = \pi c = E\Big[\sum p_i(\theta)\Big] = \text{expected total fees collected.}$$

Agent i is told that she must pay a non-refundable participation fee of

$$p(\theta_i) = \begin{cases} 0, & \theta_i < \bar{\theta} \\ \theta_i \pi_i(\theta_i) - \displaystyle\int_{\bar{theta}}^{\theta_i} \pi_i(s) ds, & \theta_i \geq \bar{\theta} \end{cases}$$

Both $\pi_i(\theta_i)$ and $p(\theta_i)$ are nonnegative and increasing in θ_i.

When $\theta_i < \bar{\theta}$ the agent is excluded, $1_i(\theta) = 0$.

Requirement (i) of individual rationality is respected since

$$\theta_i \pi_i(\theta_i) - p(\theta_i) \geq 0$$

and incentive compatibility is respected since

$$p'(t) = t\pi_i'(t)$$
$$\frac{d}{dt}\left[\theta_i \pi_i(t) - p(t)\right] = (\theta_i - t)\pi_i'(t) = 0 \implies t = \theta_i.$$

6 Optimal design

The optimal design is to provide service iff

$$1(\theta) = 1 \iff \sum_i \Big(\theta_i + \lambda(2\theta_i - 1)\Big)^+ \geq (1+\lambda)c$$

and then allow agent i to access it iff

$$\theta_i + \lambda(2\theta_i - 1) \geq 0$$

equivalently $\theta_i \geq \bar{\theta} = \lambda/(1+2\lambda)$, i.e.

$$\text{agent } i \text{ admitted, } 1_i(\theta) = 1 \iff 1(\theta) = 1 \text{ and } \theta_i \geq \bar{\theta}.$$

The parameter λ is chosen so that the expected fees equal the expected cost.

- Notice that the upfront fee of $p(\theta_i)$ is collected whether or not agent i receives service. So we are relaxing condition (iii) that excluded, those agents are neither compensated nor charged any fee (no ex-post regrets for excluded agents). This can be unrealistic.

  The expected payment of agent i is

$$E[p(\theta_i)] = \int_{\bar{\theta}}^{1} \left[\theta_i \pi_i(\theta_i) - \int_{\bar{\theta}}^{\theta_i} \pi_i(s)ds\right] d\theta_i$$
$$= E[\pi_i(\theta_i)(2\theta_i - 1)] = E[1_i(\theta)(2\theta_i - 1)].$$

- Ex-ante budget balance is satisfied:

$$E\left[\sum_i p(\theta_i)\right] = E\left[1(\theta)\right]c = \pi c,$$

  but not the ex-post form (ii), for which we would need

$$\sum_i 1_i(\theta)p_i(\theta) = 1(\theta)c.$$

- As c increases we need increasing λ. The (expected) welfare is

$$E\left[\sum_i 1_i(\theta)\theta_i - 1(\theta)c\right] = E\left[\sum_i \pi_i(\theta_i)\theta_i - \pi(\theta)c\right]$$

  The ex-ante budget balance constraint is

$$E\left[\sum_i p(\theta_i) - 1(\theta)c\right] = E\left[\sum_i 1_i(\theta_i)(2\theta_i - 1) - 1(\theta)c\right] \geq 0$$

  and this leads to the Lagrangian formulation

$$\underset{\substack{1_1(\cdot),\ldots,1_n(\cdot),1(\cdot)\\ 0\leq 1_i(\cdot)\leq 1(\cdot)\leq 1,\forall i}}{\text{maximize}} \quad E\left[\sum_i 1_i(\theta)\Big(\theta_i + \lambda(2\theta_i - 1)\Big) - (1+\lambda)1(\theta)c\right]$$

  The maximized welfare is then

$$\min_\lambda E\left[\sum_i \left(\Big(\theta_i + \lambda(2\theta_i - 1)\Big)^+ - (1+\lambda)c\right)^+\right].$$

  Social welfare decreases with c.

  $\lambda = 0$ is optimal for small c.

  $\lambda = \infty$ is optimal for maximal realistic c.

7 Construction of (ii): ex-post budget balance

7.1 Canonical solution

[1] P. Crampton, R. Gibbons, and P. Klemperer. Dissolving a partnership efficiently, *Econometrica*, vol. 55, pp. 615–632, 1987.

[2] P. Norman and T. Borgers. A note on cost-covering under interim participation constraints: The case of independent types, 2005.

For each θ let the **ex-post budget imbalance** (underpayment of cost) be

$$x(\theta) = 1(\theta)c - \sum_i p_i(\theta).$$

Ex-ante budget balance implies $E[x(\theta)] = 0$.

Pick two agents, say 1 and 2. Let new fees be as follows

$$p_1^*(\theta) = p_1(\theta_1) + x(\theta) - E\big[x(\theta) \mid \theta_1\big]$$

$$p_2^*(\theta) = p_2(\theta_2) + E\big[x(\theta) \mid \theta_1\big]$$

$$p_i^*(\theta) = p_i(\theta_i), \quad i \neq 1, 2.$$

So now we have ex-post budget balance:

$$\sum_i p_i^*(\theta) = \sum_i p_i(\theta) + x(\theta) = 1(\theta)c.$$

Also, agents' expected payments remain the same:

$$E\Big[p_1^*(\theta) \,\Big|\, \theta_1\Big] = E\Big[p_1(\theta_1) + x(\theta) - E[x(\theta) \mid \theta_1] \,\Big|\, \theta_1\Big] = p_1(\theta_1)$$

$$E\Big[p_2^*(\theta) \,\Big|\, \theta_2\Big] = E\Big[p_2(\theta_2) + E\big[x(\theta) \mid \theta_1\big] \,\Big|\, \theta_2\Big] = p_2(\theta_2)$$

so there is no change to the individual rationality and incentive compatibility calculations that agents will make. Thus (i) still holds. They will declare truthfully.

Unfortunately, after applying this modification requirement (iii) is certainly not respected.

$p_1^*(\theta)$ may be non-zero for $\theta_1 < \bar{\theta}$ (i.e. excluded agent pays.)

$p_2^*(\theta)$ may be non-zero for $\theta_2 < \bar{\theta}$

$p_i^*(\theta)$ may be non-zero when $\sum_i \Big(\theta_i + \lambda(2\theta_i - 1)\Big)^+ < (1 + \lambda)c$

(i.e. agents are asked to pay when service is not provided.)

There are alternative ways to create ex-post budget balance from a given ex-ante budget balanced design, but (so far as I know) none that respect (iii) (no ex-post regrets).

8 Construction of (iii): no ex-post regrets for excluded agents

8.1 New idea

It is well-known that it is impossible simultaneously to have ex-post budget balance and also full ex-post individual rationality and incentive compatibility, i.e. for all θ

$$p_1^*(\theta) \text{ may be non-zero for } \theta_1 < \bar{\theta}$$

$$p_2^*(\theta) \text{ may be non-zero for } \theta_2 < \bar{\theta}$$

$$p_i^*(\theta) \text{ may be non-zero when } \sum_i \Big(\theta_i + \lambda(2\theta_i - 1)\Big)^+ < (1+\lambda)c$$

In an ideal world we would like everything to hold ex-post:

$$\theta_i 1_i(\theta) - p_i(\theta) \geq 0$$

$$\theta_i 1_i(\theta_i, \theta_{-i}) - p_i(\theta_i, \theta_{-i}) \geq \theta_i 1(t_i, \theta_{-i}) - p_i(t_i, \theta_{-i}), \quad \text{for all } t_i$$

$(t_i, \theta_{-i}) = (\theta_1, \dots, t_i, \dots, \theta_n)$.

But something must give!

Can we do something modest, but still worthwhile? A minimal extra requirement would be:

$$1_i(\theta) = 0 \implies p_i(\theta) = 0.$$

So although (iii) does not hold for all agents, at least it does for the excluded ones. Call this **no ex-post regrets for excluded agents**.

8.2 non, je ne regrette rien

9 A warm up: a different construction of ex-post budget balance

$$c(\theta) = \begin{cases} 0 & \text{service not provided} \\ c & \text{service provided} \end{cases}$$

Suppose we start with a mechanism that is ex-ante budget balanced.

Now think of θ (vector of private values) taking just a finite number of discrete values, all equally likely.

That is, θ_i can take, with equal likelihood, the values $t_1 < t_2 < \cdots < t_m$

Perhaps private values are uniformly distributed over integers 1 to m. So $N = m^n$ different θ.

If it is not ex-post budget balanced then there is some $\theta = (\theta_1, \dots, \theta_n)$ for which

$$p_1(\theta) + \cdots + p_n(\theta) > c(\theta).$$

There must exist ϕ for which

$$p_1(\phi) + \cdots + p_n(\phi) < c(\phi).$$

Figure 1: *'non, je ne regrette rien'* (Edith Piaf)

- Case I.

  If there is some i such that $\theta_i = \phi_i$.

  Then adopt new payments:

$$\begin{aligned} p_i^*(\theta) &= p_i(\theta) - \epsilon \\ p_i^*(\phi) &= p_i(\phi) + \epsilon \end{aligned}$$

  Increase ϵ until there is exact budget balance at either θ, ϕ (or both).

  - The outcomes θ and ϕ are equally likely.
  - Incentive compatibility conditions do not change because

$$\begin{aligned} &E\Big[p_i^*(\theta) \;\Big|\; \theta_i\Big] - E\Big[p_i(\theta) \;\Big|\; \theta_i\Big] \\ &\quad = \frac{1}{N}\Big[\Big(p_i(\theta) - \epsilon\Big) + \Big(p_i(\phi) + \epsilon\Big) - \Big(p_i(\theta) + p_i(\phi)\Big)\Big] \\ &\quad = 0. \end{aligned}$$

- Case II.

  If $\theta_1 \neq \phi_1$ and $\theta_2 \neq \phi_2$.

  Then adopt new payments:

$$\begin{aligned} p_1^*(\theta_1, \theta_2, \theta_3, \ldots, \theta_n) &= p_1(\theta_1, \theta_2, \theta_3, \ldots, \theta_n) - \epsilon \\ p_1^*(\theta_1, \phi_2, \theta_3, \ldots, \theta_n) &= p_1(\theta_1, \phi_2, \theta_3, \ldots, \theta_n) + \epsilon \\ p_2^*(\theta_1, \phi_2, \theta_3, \ldots, \theta_n) &= p_2(\theta_1, \phi_2, \theta_3, \ldots, \theta_n) - \epsilon \\ p_2^*(\phi_1, \phi_2, \phi_3, \ldots, \phi_n) &= p_2(\phi_1, \phi_2, \phi_3, \ldots, \phi_n) + \epsilon \end{aligned}$$

and $p_i^* = p_i$ elsewhere. This can be represented as

	p_1	p_2
$\theta_1,\ \theta_2, \theta_3, \dots, \theta_m$	$-$	
$\theta_1, \phi_2, \theta_3, \dots, \theta_m$	$+$	$-$
$\phi_1, \phi_2, \phi_3, \dots, \phi_m$		$+$

The three 'states' above represent equally likely outcomes. Notice that incentives for agents 1 and 2 are unchanged. For example,

$$\begin{aligned} & E\Big[p_1^*(\theta)\ \Big|\ \theta_1\Big] - E\Big[p_1(\theta)\ \Big|\ \theta_1\Big] \\ &= \frac{1}{N}\Bigg[\big(p_1(\theta_1, \theta_2, \theta_3, \dots, \theta_n) - \epsilon\big) + \big(p_1(\theta_1, \phi_2, \theta_3, \dots, \theta_n) + \epsilon\big) \\ &\qquad - p_1(\theta_1, \theta_2, \theta_3, \dots, \theta_n) - p_1(\theta_1, \phi_2, \theta_3, \dots, \theta_n)\Bigg] \\ &= 0. \end{aligned}$$

As before, increase ϵ until there is budget balance in either θ or ϕ.

10 No ex-post regrets for excluded agents

Continue to assume now a discrete uniform model in which θ_i can take, with equal likelihood, the values $t_1 < t_2 < \cdots < t_m$, and $\bar{\theta} = t_h$ (some $h < m$). So agent i is excluded if $\theta_i \in \{t_1, \dots, t_{h-1}\}$.

We should consider the possibility that the mechanism now is sometimes randomizing, taking $0 < 1_i(\theta) < 1$, so to speak. But park this issue for now.

10.0.1 Step 0

Construct optional payments which satisfy ex-post budget balance. So now

$$\sum_i p_i(\theta) = c(\theta), \quad \text{for all } \theta.$$

An agent can be excluded for two reasons.

- $\theta_i < t_h = \bar{\theta}$. In this case the agent is excluded whatever the values of the other θ_j.

  In this case we say an agent is *ex-ante self-excluded*.

- $\theta_i \geq t_h$, but all agents are excluded, because

$$\sum_i \Big(\theta_i + \lambda(2\theta_i - 1)\Big)^+ < (1+\lambda)c.$$

10.0.2 Step 1

Consider first those θ for which **all** agents are ex-ante self-excluded.

$$\Theta_0 = \{\theta : \theta_i < t_h \text{ for all } i\},$$

For these θ the system is not provided and $\sum_i p_i(\theta) = 0$.

Suppose $p_i(\theta) > 0$ and $p_j(\theta) < 0$. WLOG $i = 1$ and $j = 2$. We change payments as follows for agents 1 and 2 for four different states:

	p_1	p_2
$\theta_1, \theta_2, \theta_3, \ldots, \theta_m$	$-$	$+$
$\theta_1, t_m, \theta_3, \ldots, \theta_m$	$+$	$-$
$t_m, t_m, \theta_3, \ldots, \theta_m$	$-$	$+$
$t_m, \theta_2, \theta_3, \ldots, \theta_m$	$+$	$-$

Increase ϵ until either $p_1^*(\theta)$ or $p_2^*(\theta)$ (or both) is 0.

- $E\Big[p_i(\theta) \Bigm| \theta_i\Big]$ is unchanged for every i.
- $\sum_i p_i(\theta)$ is unchanged for every θ.
- The number of non-zero payments in θ is reduced by at least 1. (The other states in which we are making changes to the payments are not in Θ_0, as $t_m \geq t_h$.)
- Repeat until no violation remains for any $\theta \in \Theta_0$.

10.0.3 Step 2

Now consider $\theta \notin \Theta_0$.

Suppose $\theta_i < t_h$ and $p_i(\theta) \neq 0$. WLOG assume $i = 1$ and $p_1(\theta) > 0$.

Since $\theta \notin \Theta$ there exists j such that $\theta_j \geq t_h$. WLOG let $j = 2$.

Recall that $E\big[p_1(\theta \mid \theta_1)\big] = 0$, for $\theta_1 < t_h$ and this is not changed by the constructions in steps 0 and 1. So there must exist $\phi \notin \Theta_0$ such that $\phi_1 = \theta_1$ and $p_1(\phi) < 0$. Similarly, there must exist k such that $\phi_k \geq t_h$ (where possibly $k = j$).

WLOG let $k = 3$. Make the following changes for three θ:

		p_1	p_2	p_3
θ	$\theta_1, \theta_2, \theta_3, \theta_4, \ldots, \theta_n$	$-$	$+$	
ϕ	$\phi_1(=\theta_1), \phi_2, \phi_3, \phi_4, \ldots, \phi_n$	$+$		$-$
θ'	$\theta_1, \theta_2, \phi_3, \phi_4, \ldots, \phi_n$		$-$	$+$

This preserves incentives but reduces by one the number of violations of (iii) that occur for the self-excluding agents. Since both $\theta_2 \geq t_h$ and $\phi_3 \geq t_h$ we are not making changes to the payments for any agent who would be self-excluding.

At this point we have a mechanism design that is ex-post budget balanced, ex-ante individually rational and incentive compatible, and which does not take any payment from a self-excluded agent.

10.0.4 Step 3

Lastly, we consider agents who are not self-excluded, but who are excluded because although $\theta_i \geq t_h$, we have

$$\sum_i \Big(\theta_i + \lambda(2\theta_i - 1)\Big)^+ < (1+\lambda)c.$$

and so all agents are excluded.

Let θ be the lexicographically least θ for which a violation to (iii) occurs for an agent who is not self-excluded.

Lexicographical ordering:

$$\theta \prec \theta' \quad \text{if the first nonzero component of } \theta - \theta' \text{ is negative.}$$

Suppose $c(\theta) = \sum_i p_i(\theta) = 0$, but $p_1(\theta) > 0$ and $p_2(\theta) < 0$. Since all self-excluding agents already pay 0, we must have $\theta_2 \geq t_h$.

Similarly, as in step 1 we do

		p_1	p_2
θ	$\theta_1, \theta_2, \theta_3, \dots, \theta_m$	$-$	$+$
θ'	$\theta_1, t_m, \theta_3, \dots, \theta_m$	$+$	$-$
θ''	$t_m, \theta_2, \theta_3, \dots, \theta_m$	$+$	$-$
θ''''	$t_m, t_m, \theta_3, \dots, \theta_m$	$-$	$+$

Repeat until $p_i(\theta) = 0$ for all i.

This need not decrease the number of violations of (iii), but it has the effect of ensuring that the lexicographically least θ for which there is a violation increases and incentives are unchanged.

Notice that no changes are made to payments of agents that are self-excluding.

So after a finite number of applications of Step 3 we will have a mechanism as desired:

- ex-ante individually rational;
- ex-ante incentive compatible;
- ex-post budget balance;
- no ex-post regrets for excluded agents.

11 Large systems

There is an extremely simple design which is asymptotically optimal as the number of agents becomes large.

Simply charge a fixed club fee and let each agent decided whether or not to pay it.

Since the number of potential participants is large, no single agent's decision (or θ_i) can have significant affect on whether the service is provided or not.

We can say that about $(1-\bar{\theta})n$ agents would be willing to pay a fee of $\bar{\theta}$. So if $(1-\bar{\theta})n \times \bar{\theta} > c$ we should be able to attract sufficient fee income to provide the service.

Having provided the service we obtain ex-post budget balance by dividing its cost equally between the agents who pre-commit to paying a share.

By calculation using the Law of Large Numbers.

Theorem: *As $n \to \infty$ this can achieves welfare per agent that is with $1/n$ of that which could be produced by an optimal design.*

12 Digression: puzzle

You are faced with a problem in which you need to make a decision.

For example, you might be choosing u_0.

You can choose $u_0 = 0$ or $u_0 = 1$. One of these is optimal.

There is something "A" which you do not know. It could be either true or false.

- If A is true then $u_0 = 1$ is definitely optimal (i.e. better than $u_0 = 0$).
- If A is false then $u_0 = 1$ is also definitely optimal.

Can you therefore conclude that $u_0 = 1$ is optimal?

13 Answer to the puzzle

Suppose we are faced with a stopping problem in which $u_0 = 1$ means "continue" and $u_0 = 0$ means "stop". We should consider the possibility that by knowing that A is true, or by knowing A is false, it is optimal to continue. But if we do not know whether A is true or false then it is optimal to stop.

It is not completely obvious that this can happen. So let's see how to construct an example in which it does. Consider a version of the secretary problem in which we will see 3 candidates. They are "best", "middle" and 'worst" (B,M,W). Of course we can only compare them with one another, so if they are presented in the order M,B,W, we would see this as $(x_1, x_2, x_3) = (1, 1, 0)$, where $x_i = 1$ if the ith candidate is the best so far. Suppose that the candidates will be presented in one of the of the following four orders, with consequent sequence of 0s and 1s, with the probabilities given:

B,M,W (1,0,0) with probability 0.2

M,B,W (1,1,0) with probability 0.3

B,W,M (1,0,0) with probability 0.2

W,M,B (1,1,1) with probability 0.3

It is a bit artificial that all 6 possible orders are not equally likely, but for the purpose of this discussion we are not trying to think about a really practical problem.

We wish to maximize the probability of stopping on the last 1.

Now suppose the unknown "A" is whether "the worst candidate will be seen last".

If "A" is true it will be 1 or 2, but if A is false it will be 3 or 4.

You can now verify that given A is true it is optimal not to stop on the first 1 ($u_0 = 1 =$ "continue"). The same decision is optimal if A is false. In both cases we win with probability 0.6.

However, if we do not know whether A is true or false, then if is optimal to stop on the first 1 ($u_0 = 0$, and win with probability 0.4). If we do not stop we will with probability 0.6 reach 1,1 and from there can win with probability of only 0.5. So our win probability is only 0.3.

14 Last arrivals problem

The question I have been thinking about concerns an application of Bruss's sum-the-odds algorithm in circumstances that we do not know the odds.

Theorem. If $X_1, \dots, X_n$ are independent random variables, with $X_i = 1$ or 0 with probabilities p_i and $q_i = 1 - p_i$, then we maximize the probability of stopping on the last 1 if we stop at the first 1 that we find amongst $X_s, X_{s+1}, \dots, X_n$, where s is the greatest integer i such that

$$\frac{p_i}{q_i} + \cdots + \frac{p_n}{q_n} \geq 1.$$

Now suppose we do not know the actual values of $p_1, \dots, p_n$. But we have seen $X_i = 1$ and somehow know that $p_i/q_i + \cdots + p_n/q_n = 1$. Can we say that it is optimal to stop? Since sum-of-odds = 1 is the borderline case in the odds theorem is it also optimal not to stop?

Let me take this a little further. Suppose that $p_i/q_i + \cdots + p_n/q_n > 1$, but there is some uncertainty about the order in which we will see the remaining variables. If A is true we will see them in order $i, i+1, \dots, n$. But if A is not true we will see them in the reverse order. Suppose that we know that $p_i/q_i + \cdots + p_n/q_n > 1$, and have just observed $X_{i-1} = 1$. Is it optimal to continue?

Notice that the value of $p_i/q_i + \cdots + p_n/q_n$ is the same when A is true as when A is false.

It would be optimal to continue if we knew A is true (since $i - 1 < s$).

For the same reason it would be optimal to continue if we knew A is false.

But what if we do not know whether A is true or false? Our difficulty is that if we do continue, then we must stop optimally thereafter; this is easy to do if we know whether A is true or false (we use the odds algorithm), but not easy if we do not know (the odds algorithm can't help us). So it can be better to stop on $X_{i-1} = 1$.