Prophet inequalities and posted price mechanisms

José Correa
Universidad de Chile

43rd Conference on the Mathematics of Operations Research

January 2018

Outline

Outline

An old basic problem

- Arthur Cayley 1875

> 4528. (Proposed by Professor Cayley) A lottery is arranged as follows: There are n tickets representing $a, b, c, \cdots$ pounds respectively. A person draws once; looks at his ticket; and if he pleases, draws again (out of the remaining $n-1$ tickets); and so on, drawing in all not more than k times; and he receives the value of the last ticket drawn. Supposing that he regulates his drawings in the manner most advantageous to him according to the theory of probabilities, what is the value of his expectation?

- Moser 1956 limit version. Consider n iid random variables $X_1, \ldots, X_n$. Sequentially look at the realizations until you decide to stop.
- Problem easily solved by dynamic programming. Explicit rules known for specific distributions.

The Prophet Inequality (PI)

- Rather than looking at the optimal stopping rule, Krengel & Sucheston 1977 ask the "Prophet vs Gambler" question

> We have n independent and nonnegative random variables $X_1, \ldots, X_n$ with $X_i \sim F_i$.
> They arrive sequentially and upon arrival reveal their value.
> Gambler with distributional knowledge, either keeps current value, or drops it and continue.
> Prophet sees the entire realization in advance and picks the maximum.

How well can the Gambler do?

- Related to the Secretary problem: n arbitrary values arrive in random order and we want to pick the maximum.

The Prophet Inequality (PI)

- How well can the Gambler do?

$$\underbrace{\sup_{t \text{ stopping time}} \mathbb{E}[X_t]}_{\text{Gambler}} \geq \frac{1}{2} \underbrace{\mathbb{E}\left[\max_{1\leq i\leq n} X_i\right]}_{\text{Prophet}}$$

Gambler vs Prophet

And $1/2$ is best possible. Krengel, Sucheston, Garling 1977

- For the Secretary problem you have to scan a fraction $1/e$ of the values and then pick the first value above the maximum seen so far.

$$\mathbb{P}(\text{pick the maximum}) = 1/e$$

- Yet another, the Prophet Secretary problem. Same as prophet inequality but r.v. arrive in random order.
Can improve the $1/2$ to $1 - 1/e \approx 0.63$.

Auctions

AUCTIONS:

- We have one item and a set N of potential buyers.
- Buyers have independent random valuations for the item.
- Buyer i valuation is $v_i \sim F_i$.
- Revenue maximizing auction: Myerson, MOR 81
  Difficult to implement and complicated.

POSTED PRICE MECHANISMS (PPM):

- Same setting
- Now customers arrive in some order (selected by the seller, or random, or worst case)
- Seller sequentially makes take-it-or-leave-it offers
- Simple, no strategic behavior, easy to implement

GOAL: compare the revenue of these mechanisms

Outline

PPM vs PI

- We are given n independent and nonnegative random variables $X_1, \dots, X_n$ with $X_i \sim F_i$.
- The rv's arrive sequentially and upon arrival reveal their value.
- PI: you pick thresholds $T_1, \dots, T_n$
- PPM: you pick prices $p_1, \dots, p_n$
- PI: if realization $x_i \geq T_i$ you stop and get x_i.
- PPM: if realization $x_i \geq p_i$ you stop and get p_i.
- PI: compare revenue against $\mathbb{E}(\max_i X_i)$.
- PPM: compare revenue against Optimal mechanism.

Myerson's Lemma: Optimal revenue equals $\mathbb{E}(\max_i c_i(X_i))$, where $c_i(x) = x - \frac{1-F_i(x)}{f_i(x)}$ is the virtual valuation function (assume monotone and nonnegative)

PPM and PI are Equivalent problems

- For any sequence of r.v. there exist thresholds achieving an expected reward of α times the expected maximum, if and only if for any sequence of r.v. there exists a PPM that achieves a revenue of α times the optimal auction.
- Key: Move to the virtual valuations and back.

Example: Two buyers $V_1, V_2 \sim U[1,2]$.

- Optimal Auction is a SPA, revenue is $\mathbb{E}(\min(V_1, V_2)) = 4/3$.
- PPM: Say set price $P = 3/2$, obtain $\frac{3}{2} \times \frac{3}{4} = \frac{9}{8}$.
- Virtual Values: $X_i = V_i - \frac{1-F_i(V_i)}{f_i(V_i)} = 2(V_i - 1) \sim U[0,2]$.
- $\mathbb{E}(\max(X_1, X_2)) = 4/3$
- Threshold: set T so that prob. stop stays at $1/2$, then $T = 1$.
- Obtain $\frac{1}{2}\mathbb{E}(X_1 | X_1 > 1) + \frac{1}{4}\mathbb{E}(X_2 | X_2 > 1) = \frac{9}{8}$

PPM and PI are Equivalent problems

Chawla et al 2010, C. Foncea, Pizarro, Verdugo 2017

- For any sequence of r.v. there exist thresholds achieving an expected reward of α times the expected maximum, if and only if for any sequence of r.v. there exists a PPM that achieves a revenue of α times the optimal auction.
- Key: Move to the virtual valuations and back.
- Fact: Consider a r.v. $V \sim F$ and let the r.v. X be the virtual valuation of V.

$$X = c(V) = V - \frac{1 - F(V)}{f(V)}.$$

  - If $q = \mathbb{P}(X \geq T)$ then $X \geq T$ iff $V \geq F^{-1}(1-q)$
  - Also $\mathbb{E}(X \mid X \geq T) = F^{-1}(1-q)$

$$\begin{aligned}\int_{F^{-1}(1-q)}^{\infty} c(v) f(v) dv &= \int_{F^{-1}(1-q)}^{\infty} v f(v) dv - \left(v(1-F(v))\big|_{F^{-1}(1-q)}^{\infty} + \int_{F^{-1}(1-q)}^{\infty} v f(v) dv \right) \\ &= q F^{-1}(1-q)\end{aligned}$$

From PI to PPM

- For $i = 1, \ldots, n$ let $V_i \sim F_i$ and X_i its virtual valuation.
- Take the PI thresholds T_i (for X_i) and $q_i = \mathbb{P}\left(X_i \geq T_i\right)$.
- Let r be the index of the first r.v. with virtual value above the threshold. Then

$$
\begin{aligned}
\mathbb{E}\left(X_r\right) &= \sum_{i=1}^{n} \mathbb{E}\left(X_i \mid i = r\right) \mathbb{P}\left(i = r\right) \\
&= \sum_{i=1}^{n} \mathbb{E}\left(X_i \mid X_i \geq T_i\right) \mathbb{P}\left(X_i \geq T_i,\ i \text{ is the first accepting}\right) \\
&= \sum_{i=1}^{n} F_i^{-1}\left(1 - q_i\right) \mathbb{P}\left(V_i \geq F_i^{-1}\left(1 - q_i\right),\ i \text{ is the first accepting}\right) \\
&= \text{ revenue of a PPM with prices} \\
&\qquad p_i = F_i^{-1}\left(1 - q_i\right) \text{ for r.v. } V_1, \ldots, V_n
\end{aligned}
$$

From PI to PPM

- Then revenue of threshold rule over the virtual valuations equals revenue of the PPM over the original valuations.
- Also expectation of maximum of $X_1, \ldots, X_n$ equals optimal auction revenue. Myerson, MOR 81
- We conclude since

$$
\begin{aligned}
\mathbb{E}\left(\text{PPM over } V_1, \ldots, V_n\right) &\geq \mathbb{E}\left(X_r\right) \\
&\geq \alpha \mathbb{E}(\max_i X_i) \\
&= \alpha \mathbb{E}(\text{rev opt auction})
\end{aligned}
$$

From PPM to PI

Converse is more involved

- Key Lemma: For any r.v. X there exist another random variable Y whose (ironed) virtual valuation is distributed as X.
- Consider $X_1, \dots, X_n$ and the corresponding $Y_1, \dots, Y_n$.
- Take a PPM with a guarantee of α for sequence $Y_1, \dots, Y_n$ (w.r.t max virtual valuation)
- Transform the prices into thresholds
- Resulting thresholds achieve a guarantee of α for $X_1, \dots, X_n$ (w.r.t max)

Outline

The Prophet Inequality: Simple proof

- Many proofs Garling 78, Hill & Kertz 81, Samuel-Cahn 83, etc
- Recent one based on Kleinberg Weinberg 2012
- Pick a threshold T and accept first value above T (Anonymous!!)
- Let r be the index of the first r.v. above the threshold
- $p = \mathbb{P}(\max X_i > T)$

Note that

$$\mathbb{P}(X_r > x) \geq \begin{cases} p & x \leq T \\ (1-p)\mathbb{P}(\max X_i > x) & x > T \end{cases}$$

Indeed,

$$\begin{aligned} \mathbb{P}(X_r > x) &= \sum_{i=1}^{n} \mathbb{P}(X_i > x) \prod_{j<i} \mathbb{P}(X_j \leq T) \\ &\geq (1-p) \sum_{i=1}^{n} \mathbb{P}(X_i > x) \geq (1-p)\mathbb{P}(\max X_i > x) \end{aligned}$$

The Prophet Inequality: Simple proof

$$\text{Gambler} = \mathbb{E}(X_r) = \int_0^\infty \mathbb{P}(X_r > x)$$

$$\geq \int_0^T p dx + \int_T^\infty (1-p)\mathbb{P}(\max X_i > x) dx$$

$$\text{Note that } \mathbb{E}(\max X_i) = \int_0^\infty \mathbb{P}(\max X_i > x) dx \leq T + \int_T^\infty \mathbb{P}(\max X_i > x) dx$$

$$\geq pT + (1-p)\left(\mathbb{E}(\max X_i) - T\right)$$

Pick T s.t. $\begin{cases} T = \frac{\mathbb{E}(\max X_i)}{2} \\ p = \frac{1}{2} \end{cases}$ to obtain Gambler $\geq \frac{1}{2}$Prophet.

Tightness

- Worst case: Two rv's. $X_1 = 1$ a.s., and

  $$X_2 = \begin{cases} 1/\varepsilon & \text{w.p. } \varepsilon \\ 0 & \text{w.p. } 1-\varepsilon \end{cases}$$

  Gambler gets 1 while Prophet gets $\varepsilon(1/\varepsilon) + (1-\varepsilon) \approx 2$.
- The constant $1/2$ cannot be beaten even if we choose non-anonymous thresholds.
- So single threshold strategies are optimal!!

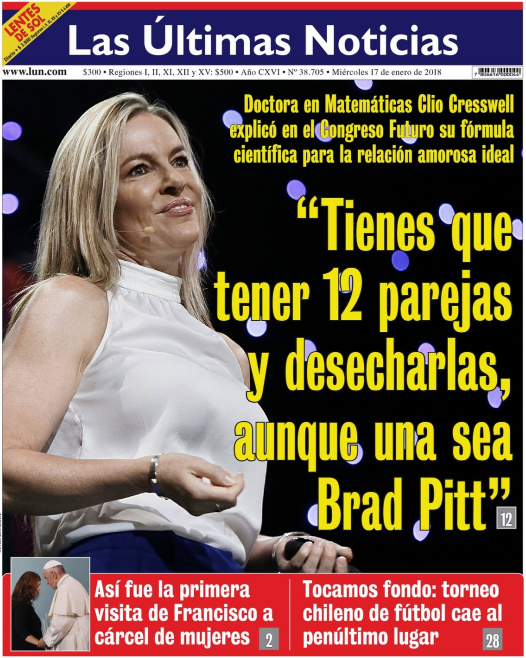
Las Últimas Noticias

www.lun.com $300 • Regiones I, II, XI, XII y XV: $500 • Año CXVI • Nº 38.705 • Miércoles 17 de enero de 2018

Doctora en Matemáticas Clio Cresswell explicó en el Congreso Futuro su fórmula científica para la relación amorosa ideal

"Tienes que tener 12 parejas y desecharlas, aunque una sea Brad Pitt" 12

Así fue la primera visita de Francisco a cárcel de mujeres 2

Tocamos fondo: torneo chileno de fútbol cae al penúltimo lugar 28

"You need to have 12 couples and discard all of them, even if one is Brad Pitt"

"...then you keep the first better"

Recap

- Prophet inequality: Basic problem in optimal stopping.
- Equivalent to Posted price mechanisms through virtual values
- Basic setting: $X_1, \ldots, X_n$ nonnegative r.v.s.
  - Compare gambler, i.e., good stopping rule or algorithm to stop
  - with Prophet, i.e., expectation of the maximum.
- Classic PI: Gambler can do $1/2$ of the prophet by pick a threshold and accept the first r.v. above it.
- The constant $1/2$ cannot be beaten even with non-anonymous thresholds.
- So single threshold strategies are optimal!!

Outline

Personalization and Adaptivity

- For the basic PI single threshold strategies are optimal
  - $T = \mathbb{E}(\max X_i)/2$ or
  - T such that $\mathbb{P}(\max X_i > T) = 1/2$

What if the order can be selected by the Gambler or the order is RANDOM (Prophet Secretary)?

- Optimal strategies through exponential dynamic programs. Complexity open.
- Role of Personalization?
  Thresholds may depend on the index of the r.v. that is seen.
- Role of Adaptivity?
  Strategy may be adjusted depending on what is left.

Observations

- What if the order can be selected by the Gambler?
  Can obtain a fraction $1 - 1/e$ Chawla et al., 2010, Yan 2011
  Uses personalization
- What if the order is RANDOM? Prophet Secretary
  Can obtain a fraction $1 - 1/e$ Esfandiari et al. 2015
  Uses adaptivity
  Can obtain a fraction $1 - 1/e$ nonadaptively C. et al. 2017
  Uses personalization
  $1 - 1/e$ is best possible nonadaptively C. et al. 2017
  $1 - 1/e + 1/400$ just obtained Azar et al. 2018
  Uses personalization and adaptivity
- Prophet Secretary, single threshold? Ehsani et al. 2018
- What about iid rv's as in Cayley-Moser? C. et al. 2017

Single Threshold: Is $1/2$ best possible?

- Consider n random variables
- $n - 1$ are deterministic and always give 1.
- The other gives n with probability $1/n$ and zero with probability $1 - 1/n$.
- The r.v. arrive in random order.

$$\mathbb{E}(\max X_i) = n \times 1/n + 1 \times (1 - 1/n) \approx 2$$

Now fix a threshold T.

$\rightarrow$ If $T < 1$, Gambler gets n w.p. $\frac{1}{n^2}$, and 1 ow. So Gambler ≈ 1
$\rightarrow$ If $T \geq 1$ Gambler gets n w.p. $\frac{1}{n}$. So Gambler $= 1$

- So $1/2$ is best possible.

Single random threshold

- To beat $1/2$ one can use personalization or adaptivity.
- Not needed. A random threshold does it! Ehsani et al. 2018
- Take the example again. Set $T = 1$ and break ties at random.
- Say you accept a 1 w.p. $1/n$.
  Equivalently think that $X_1, \ldots, X_{n-1}$ are 1 w.p. $1/n$ and 0 o.w. and that you accept the first nonzero.

$$\begin{aligned} \mathbb{P}(\text{Gambler gets something}) &= 1 - \mathbb{P}(\text{Gambler gets nothing}) \\ &= 1 - (1 - 1/n)^n \approx 1 - 1/e \\ \mathbb{E}(\text{Gambler} | \text{ gets something}) &= n \times 1/n + 1 \times (1 - 1/n) \approx 2. \end{aligned}$$

Therefore $\mathbb{E}(\text{Gambler}) \approx 2(1 - 1/e) \approx (1 - 1/e)\mathbb{E}(\max X_i)$.

Single random threshold

- Can get $1 - 1/e$ in general!
- Here is a *simple* proof C. Cristi, Saona, last week
- Proof for continuous and strictly increasing distributions.
- Let $X_1, \ldots, X_n$ be nonnegative r.v. $X_i \sim F_i$, they come in random order.
- Set a threshold T so that $\mathbb{P}(\max X_i < T) = \prod F_i(T) = 1/e$.
- Stop at (a random) time r; the first time you see T or more.
- THM: $\mathbb{E}(X_r) \geq (1 - 1/e)\mathbb{E}(\max X_i)$.

Proof

- We prove a stronger statement: $\mathbb{P}(X_r > x) \geq (1 - 1/e)\mathbb{P}(\max X_i > x)$
- σ random permutation. So X_i comes at time $\sigma(i)$.

$$\begin{aligned}
\mathbb{P}(r = \sigma(i) | X_i > T) &= \sum_{S \subset [n] \setminus i} \frac{1}{|S| + 1} \prod_{j \in S} (1 - F_j(T)) \prod_{j \notin S, j \neq i} F_j(T) \\
&= \frac{1}{eF_i(T)} \sum_{S \subset [n] \setminus i} \frac{1}{|S| + 1} \prod_{j \in S} \frac{1 - F_j(T)}{F_j(T)} \\
&\geq \frac{1}{eF_i(T)} \cdot \min_{\prod_j x_j = \frac{1}{eF_i(T)}} \sum_{S \subset [n] \setminus i} \frac{1}{|S| + 1} \prod_{j \in S} \frac{1 - x_j}{x_j} \\
&\color{red}{\geq 1 - \frac{1}{e}}
\end{aligned}$$

For the minimization problem we change variables $y_j = -\ln(x_j)$ and note that the resulting objective is Schur-Convex i.e., $(y_i - y_j)(\partial g / \partial y_i - \partial g / \partial y_j) \geq 0$ so by the Schur-Ostrowski criterion it is minimized when all variables are equal.

Proof

- We prove a stronger statement: $\mathbb{P}(X_r > x) \geq (1 - 1/e)\mathbb{P}(\max X_i > x)$
- To finish the proof we condition on the stopping time

$$\begin{aligned}
\mathbb{P}(X_r > x) &= \sum_{i=1}^{n} \mathbb{P}(X_i > x | r = \sigma(i))\mathbb{P}(r = \sigma(i)) \\
&\geq \sum_{i=1}^{n} \mathbb{P}(X_i > x | r = \sigma(i))(1 - \frac{1}{e})\mathbb{P}(X_i > T) \\
&= (1 - \frac{1}{e}) \sum_{i=1}^{n} \frac{\mathbb{P}(X_i > x)}{\mathbb{P}(X_i > T)}\mathbb{P}(X_i > T) \\
&\geq (1 - \frac{1}{e})\mathbb{P}(\max X_i > x)
\end{aligned}$$

Personalization

Non-Adaptive Setting

C., Foncea, Hoeksma, Oosterwijk, Vredeveld 2017

- Think of a direct mail campaign.
  Buyers are simultaneously contacted by email with a personalized offer. They respond in random order and the offer is valid until the item is sold.
- $X_1, \ldots, X_n$ arrive in random order.
- Gambler precomputes thresholds $T_1, \ldots, T_n$.
- Whenever X_i comes stop iff $x_i \geq T_i$.
- Thresholds ONLY depend on a priori knowledge of the r.v's.
- Cannot beat single threshold bound even with iid r.v.!!

Personalization

Non-Adaptive Setting

- IID Instance where no non-adaptive algorithm can achieve a factor better than $1 - 1/e$.
- Consider n^2 r.v.'s with i.i.d. valuations

$$X = \begin{cases} \frac{n}{e-2} & \text{w.p. } \frac{1}{n^3}\,, \\ 1 & \text{w.p. } \frac{1}{n}\,, \\ 0 & \text{w.p. } 1 - \frac{1}{n} - \frac{1}{n^3}\,. \end{cases}$$

- Expectation of maximum goes to $\frac{e-1}{e-2}$ as $n \to \infty$.
- BEST threshold strategy: Set threshold 1 for n r.v.'s and threshold $\frac{n}{e-2}$ for the rest.
- Revenue approaches $\frac{(e-1)^2}{e(e-2)}$, as $n \to \infty$.

Adaptivity

- Think of business class upgrades at check-in. Customers arrive in random order (say at times $1, 2, \ldots, n$) and are offered a price for a seat upgrade.
- If Customer i arrives at time k is offered a price that depends on the customers that already declined.
- If she accepts, she immediately gets the item.
- Prices adapt to the current situation.

Known result: There is an adaptive strategy obtaining a fraction $1 - 1/e + 1/400$ of optimal revenue. Azar et al 2018

Known result: For IID r.v.s we can get to 0.745 and this is best possible C., Foncea, Hoeksma, Oosterwijk, Vredeveld 2017

Outline

The IID Prophet Inequality

- Initiated by Gilbert and Moser 1965
- Hill and Kertz 82 provide some recursively defined upper bounds that computationally evaluate to 0.745.
- Also prove lower bound of $1 - 1/e = 0.63\ldots$.
- Conjecture that tight bound is $1 - 1/(e+1) = 0.731$
- Samuel-Cahn 84 reports that Kertz proved that the upper bound is actually $1/\beta^* \approx 0.745$ the unique solution to

$$\int_0^1 \frac{1}{y(1-\ln(y)) + (\beta - 1)} dy = 1. \qquad (1)$$

- Kertz 86 and Saint-Mount 02 conjectured that this constitutes the best possible upper bound.
- We prove this conjecture.

C., Foncea, Hoeksma, Oosterwijk, Vredeveld 2017

Algorithm's Overview

INPUT: Random variables X_i, $i = 1, \ldots, n$, with distribution F.

ALGORITHM:

- Partition interval $[0, 1]$ into intervals $A_i = [\varepsilon_{i-1}, \varepsilon_i]$, s.t. $\varepsilon_0 = 0$, $\varepsilon_n = 1$.
- Sample q_i from A_i with density $f_i(q) = (n-1)(1-q)^{n-2}/\gamma_i$. here γ_i is the normalization.
- When the $i-$th r.v comes, stop if value at least $F^{-1}(1 - q_i)$.

NOTES:

1. Appropriate choice of ε_i's gives the bound.
2. Threshold is high for the first comers and lowers later on.
3. Can get rid of randomization.
4. Easy to extend to general distributions.

Analysis Summary

- Gambler $= \sum_{i=1}^{n} \int_{\varepsilon_{i-1}}^{\varepsilon_i} (n-1)(1-q)^{n-2} R(q) dq \cdot \rho_i$.
- Prophet $= n \int_0^1 (n-1)(1-q)^{n-2} R(q) dq$.
- TRICK: Choose intervals A_i such that $\rho_1 = \rho_2 = \ldots = \rho_n$.
- This implies that the revenue is at least $\frac{1}{n\gamma_1} OPT$.
- The rest of the proof is to bound the term $(n\gamma_1)$ by $1.341 = 1/0.745$.
- Set up a recursion whose solution determines γ_1.
- Approximate the recursion with an ordinary differential equation.

NOTE:

Bound is best possible

Hill & Kertz, *Ann. Probab.* 1982

A useful expression

- Let $X_1, \ldots, X_n$ be non-negative iid. rv's, with distribution F.
- We will assume F is continuous and increasing for simplicity.
- Let $R(q) = \int_0^q F^{-1}(1-\theta)d\theta$.
- By Fubini and integration by parts Prophet gets:

$$\begin{aligned}\mathbb{E}(\max\{X_1, \ldots, X_n\}) &= \int_0^\infty 1 - F^n(t)dt = \int_0^1 F^{-1}(\sqrt[n]{z})dz \\ &= n\int_0^1 (1-q)^{n-1}F^{-1}(1-q)dq \\ &= n\int_0^1 (n-1)(1-q)^{n-2}R(q)dq\,.\end{aligned}$$

Local reward of quantile

Suppose we face a rv and accept with probability q (i.e., stop if value above $\tau(q) = F^{-1}(1-q)$).
Then the expected reward in that step equals $R(q)$. Indeed, the reward can be calculated as:

$$\begin{aligned}
&\mathbb{P}(X > \tau(q))\mathbb{E}[X|X > \tau(q)] \\
&= \int_0^\infty \mathbb{P}(X > t, X > \tau(q))dt \\
&= \int_{\tau(q)}^\infty 1 - F(t)dt \\
&= \int_0^q F^{-1}(1-\theta)d\theta = R(q).
\end{aligned}$$

Quantile stopping rule

- We take a quantile approach.
- We define quantiles $0 < q_1 < \cdots < q_n < 1$ and stop in the i-th step if $X_i \geq \tau(q_i)$.
- To define the $q_i's$ partition the interval $A = [0, 1]$ into n intervals $A_i = [\varepsilon_{i-1}, \varepsilon_i]$, with $0 = \varepsilon_0 < \varepsilon_1 < \ldots < \varepsilon_{n-1} < \varepsilon_n = 1$.
- Draw q_i at random from A_i, according to the density function $f_i(q) = \frac{\psi(q)}{\gamma_i}$, where $\psi(q) = (n-1)(1-q)^{n-2}$ and $\gamma_i = \int_{q \in A_i} \psi(q) dq$.

Right choice of ε_i...

Our Reward

- At step i reward equals $R(q_i)$.
- Probability that we get to step i is $\prod_{j=1}^{i-1}(1-q_j)$.

By linearity of expectation and independence of q_i's, Gambler gets:

$$\text{Gambler} = \sum_{i=1}^{n} \mathbb{E}(R(q_i)) \prod_{j=1}^{i-1} \mathbb{E}(1-q_j)$$

$$= \sum_{i=1}^{n} \int_{\varepsilon_{i-1}}^{\varepsilon_i} (n-1)(1-q)^{n-2} R(q) dq \frac{\prod_{j=1}^{i-1} \int_{\varepsilon_{j-1}}^{\varepsilon_j} \psi(q)(1-q)dq}{\prod_{j=1}^{i} \gamma_i}$$

$$= \sum_{i=1}^{n} \int_{\varepsilon_{i-1}}^{\varepsilon_i} (n-1)(1-q)^{n-2} R(q) dq \cdot \rho_i.$$

Where $\rho_1 = \frac{1}{\gamma_1}$ and $\rho_{i+1} = \rho_i \frac{\int_{\varepsilon_{i-1}}^{\varepsilon_i} \psi(q)(1-q)dq}{\gamma_{i+1}}$

Choosing the ε_i's

Choose $\varepsilon_1, \ldots, \varepsilon_{n-1}$ such that $\rho_1 = \rho_2 = \ldots = \rho_n$, then

$$\text{Prophet} = \mathbb{E}(\max\{X_1, \ldots, X_n\}) = n\gamma_1 \text{Gambler}.$$

This choice amounts to choosing ε_i's such that

$$\int_{\varepsilon_i}^{\varepsilon_{i+1}} \psi(q)dq = \int_{\varepsilon_{i-1}}^{\varepsilon_i} \psi(q)(1-q)dq.$$

Since $\psi(q) = (n-1)(1-q)^{n-2}$, and substituting $x_i = 1 - \varepsilon_i$ this is equivalent to

$$\frac{{x_{i-1}}^n}{n} - \frac{{x_i}^n}{n} = \frac{{x_i}^{n-1}}{n-1} - \frac{{x_{i+1}}^{n-1}}{n-1}, \tag{2}$$

where $x_0 = 1$ and $x_n = 0$.

$$\boxed{\text{Quantity of interest } n\gamma_1 = n\int_0^{\varepsilon_1} \psi(q)\,dq = n(1 - {x_1}^{n-1})}$$

Concluding through an ODE

- Consider $y(t) : [0,1] \to \mathbb{R}$, defined by the ODE:

$$\begin{aligned} y' &= y(\ln(y) - 1) - (\beta - 1)\,, \\ y(0) &= 1\,. \end{aligned}$$

- We prove that if $n(1 - {x_1}^{n-1}) > \beta$ then ${x_i}^{n-1} < y(\frac{i}{n})$. ($y(1) := \lim_{t \uparrow 1} y(t)$ is the continuous extension of $y(t)$).
- Take β such that $y(1) = 0$ to contradict $x_n = 0$.
- $y(t)$ is invertible so look at t as a function of y. We know $t(1) = 0$ and we want to choose β such that $t(0) = 1$.

$$\begin{aligned} t(1) &= t(0) + \int_0^1 \frac{dt}{dy} dy = 1 + \int_0^1 \frac{1}{\frac{dy}{dt}} dy \\ &= 1 - \int_0^1 \frac{1}{y(1 - \ln(y)) + (\beta - 1)} dy. \end{aligned}$$

This yields $\beta^* \approx 1.3415$, and thus $n\gamma_1 \leq 1.3415$.

Final Remarks

- Results extend to sequential posted pricing context.
- OPEN: What about general rv's (independent but not identical) that arrive in random order?

  know how to obtain $0.63 + 1/400$. Can we do better? Is IID the worst case?

- Do not know how to do this EVEN if Gambler can choose the order.
- What about the IID case but when we do not know the distribution (as in the secretary problem)?