Dynamic Allocation and Pricing: A Mechanism Design Approach

Benny Moldovanu

September 2011

Modern Revenue Management

- Starts with US Airline Deregulation Act of 1978.
- Today mainstream business practice (airlines, trains, hotels, car rentals, holiday resorts, advertising, intelligent metering devices, etc..)
- Considerable gap between practitioners and academics in the field
- Major academic textbook : *The Theory and Practice of Revenue Management* by K. T. Talluri and G.J. van Ryzin

Basic RM Questions (Talluri & van Ryzin)

- Quantity decisions: How to allocate capacity/output to different segments, products or channels ? When to withhold products from the market ?
- Structural decisions: Which selling format ? (posted prices, negotiations, auctions, etc..). Which features for particular format ?(segmentation, volume discounts, bundling, etc..)
- Pricing decisions: How to set posted prices, reserve prices ? How to price differentiate ? How to price over time ? How to markdown over life time ?

Towards a Modern Theory of RM

- Necessary blend of

1. The elegant **dynamic models** from the OR, MS, CS , Econ (search) literatures with historical focus on "**grand, centralized optimization**" and/or "**ad-hoc**", intuitive mechanisms.
2. The rich, classical mechanism design literature with historical focus on **information/incentives** in **static** settings.

- Blend fruitful for numerous applications.

Reviewed Papers (joint work with A. Gershkov)

- Dynamic Revenue Maximization with Heterogeneous Objects: A Mechanism Design Approach, *AEJ: Microeconomics 2009*
- Efficient Sequential Assignment with Incomplete Information, *GEB 2010*
- Revenue Maximization for the Dynamic Knapsack Problem, *TE* 2011 (also joint with D. Dizdar)
- Learning About The Future and Dynamic Efficiency, *AER 2009*
- Optimal Search, Learning, and Implementation, *mimeo 2010.*
- Efficient Dynamic Allocation with Strategic Arrivals, *mimeo 2011*

The Gallego & van Ryzin Model (*MS 94*) I

- Revenue Maximizing (RM) seller has n identical objects that perish after deadline T.
- Agents arrive according to a Poisson process with intensity λ.
- Agents' values are private information, represented by I.I.D random variables $X_i = X$ on $[0, +\infty)$ with common c.d.f. F.
- Agents desire one object only, and can only be served upon arrival (no recall)
- After an item is assigned, it cannot be reallocated.

The Poisson Process

- The Poisson process with intensity λ is a continuous-time counting process $\{N(t), t \geq 0\}$ such that:

1. $N(0) = 0$
2. Independent increments: the numbers of events in disjoint time intervals are random variables independent from each other.
3. Stationary increments: the probability distribution of the number of events in any time interval only depends on the length of the interval.
4. $N(t+\tau) - N(t)$, the number of events in time interval $(t, t+\tau]$ follows a Poisson distribution with parameter $\lambda\tau$:

$$P[N(t+\tau) - N(t) = k] = \frac{e^{-\lambda\tau}(\lambda\tau)^k}{k!}, \ \ k = 0, 1, 2...$$

5. The probability distribution of the waiting time until the next event is an exponential distribution.

Gallego & van Ryzin II

- Gallego & van Ryzin restrict attention to posted prices: at each point in time t, seller sets price p_t that needs to be paid by any buyer that arrives at t.
- **Main Results:**
    1. Optimal posted-price revenue is concave in the number of objects, and in time until deadline.
    2. Time pattern of optimal prices: Downward trend, interrupted by upward jumps after each sale.
    3. Fixed price is approximately optimal if $\min(n, \lambda T)$ is large enough.

New Research Questions

- General Mechanisms
- Multiple, Heterogenous Objects
- Multi-Unit Demand
- Learning about Demand
- Recall and Strategic Arrivals

Derman-Liebermann-Ross (MS 72)

- each item $i = 1, .., m$ is characterized by a "quality" q_i with $0 \leq q_m \leq q_{m-1} \leq ... \leq q_1$
- n agents arrive sequentially, one agent per period
- agents can only be served upon arrival
- after an item is assigned, it cannot be reallocated
- each agent is characterized by a "type" x_j
- if an item with type $q_i \geq 0$ is assigned to an agent with type x_j, this agent enjoys a utility of $q_i x_j$
- agents' types are IID random variables $X_i = X$ on $[0, +\infty)$ with common cdf. F

DLR Model II

Theorem (DLR 72)

Consider the arrival of an agent with type x_k in period $k \geq 1$. There exist $k+1$ constants $0 = a_{k,k} \leq a_{1,k} \leq ... \leq a_{0,k} = \infty$ such that:

- *The dynamically efficient policy assigns the item with quality $q_{(i)}$ if $x_k \in (a_{i,k}, a_{i-1,k}]$*
- *The constants are given recursively by*

$$a_{i,k+1} = \int_{a_{i,k}}^{a_{i-1,k}} xdF(x) + a_{i,k}F(a_{i,k}) + a_{i-1,k}[1 - F(a_{i-1,k})]$$

  where we set $+\infty \cdot 0 = -\infty \cdot 0 = 0$.

- *In a problem with n periods total expected welfare is given by $W_n = \sum_{i=1}^{n} q_{(i)} a_{i,n+1}$.*

Majorization: Definition

Definition

For any $n-$tuple $\gamma = (\gamma_1, \gamma_2, .., \gamma_n)$ let $\gamma_{(j)}$ denote the *jth* largest coordinate (so that $\gamma_{(n)} \leq \gamma_{(n-1)} \leq ... \leq \gamma_{(1)}$). Vector $\alpha = (\alpha_1, \alpha_2, .., \alpha_n)$ is *majorized* by $\beta = (\beta_1, \beta_2, .., \beta_n)$ and we write $\alpha \prec \beta$ if the following system of $n-1$ inequalities and one equality is satisfied:

$$\begin{aligned}
\alpha_{(1)} &\leq \beta_{(1)} \\
\alpha_{(1)} + \alpha_{(2)} &\leq \beta_{(1)} + \beta_{(2)} \\
... &\leq ... \\
\alpha_{(1)} + \alpha_{(2)} + ..\alpha_{(n-1)} &\leq \beta_{(1)} + \beta_{(2)} + \beta_{(n-1)} \\
\alpha_{(1)} + \alpha_{(2)} + .. + \alpha_{(n)} &= \beta_{(1)} + \beta_{(2)} + .. + \beta_{(n)}
\end{aligned}$$

We say that α is *weakly sub-majorized* by β and we write $\alpha \prec_w \beta$ if all relations above hold with weak inequality.

Majorization: Main Results (Marshall & Olkin, 79)

Theorem

1. $\alpha \prec \beta$ *if and only if there exists a doubly stochastic matrix* Q *such that* $\alpha = \beta Q$.
2. *A function* $\Psi : \mathbb{R}^n \to \mathbb{R}$ *is called Schur-convex if* $\Psi(\alpha) \leq \Psi(\beta)$ *for any* α, β *such that* $\alpha \prec \beta$.
3. *Let* $\Psi : \mathbb{R}^n \to \mathbb{R}$ *be symmetric with continuous partial derivatives. Then* Ψ *is Schur-convex if and only if for all* $(y_1, ..., y_n) \in \mathbb{R}^n$ *and all* $i, j \in \{1, .., n\}$ *it holds that*

$$\left(\frac{\partial \Psi (y_1, ..., y_n)}{\partial y_i} - \frac{\partial \Psi (y_1, ..., y_n)}{\partial y_j} \right) (y_i - y_j) \geq 0.$$

Majorization in the DLR Model

- Let $X_{(i),k}$ denote the $i - th$ highest order statistic out of k copies of X, and let $\mu_{(i),k}$ denote its expectation.
- Let S_n denote the set of permutations of $\{1, 2, .., n\}$. For each permutation $\sigma \in S_n$, define $\mu_n^\sigma = \left(\mu_{(\sigma(1)),n}\ ,\ \mu_{(\sigma(2)),n}\ , \ldots,\ \mu_{(\sigma(n)),n}\right)$
- Total expected welfare in a static problem where n agents arrive simultaneously is given by $\sum_{i=1}^{n} q_{(i)}\mu_{(i),n}$.

Theorem

The $n-$tuple $\{a_{i,n+1}\}_{i=1}^{n}$ is majorized by the $n-$tuple $\{\mu_{i,n}\}_{i=1}^{n}$. In particular, there exist n permutations $\{\sigma^i\}_{i=1}^{n}$ and non-negative weights $\{w_i\}_{i=1}^{n}$ such that $\sum_{i=1} w_i = 1$ and $\sum_{i=1}^{n} w_i \mu_n^{\sigma^i} = (a_{1,n+1}, a_{2,n+1}, .., a_{n,n+1})$.

Albright's Model (*MS 74*) I

- Welfare Maximizing (WM) seller has m items.
- Each item $i = 1, .., n$ is characterized by a "quality" q_i with

$$0 \leq q_n \leq q_{n-1} \leq ... \leq q_1$$

- If an item with quality $q_i \geq 0$ is assigned to an agent with type x_j this agent has utility $q_i x_j$
- Complete information: current type is known upon arrival; future types are IID random variables $X_i = X$ on $[0, +\infty)$ with common c.d.f. F.
- Poisson arrivals, unit demand, deadline, etc..as in Gallego & van Ryzin.

The Welfare Maximizing (WM) Allocation

Theorem (Albright)

Denote by Π_t the set of items available at t, with cardinality k_t. There exist $n+1$ unique functions

$$\infty \equiv y_0(t) \geq y_1(t) \geq ..y_n(t) \geq 0, \ \forall t$$

which do not depend on the q's such that if an agent with type x arrives at time t, it is optimal to assign him the j'th highest element of Π_t if $x \in [y_j(t), \ y_{j-1}(t))$ and not to assign any object if $x < y_{k_t}(t)$.

1. *For each k, the function $y_k(t)$ satisfies* $y_k'(t) = -\lambda \int_{y_k}^{y_{k-1}} (1 - F(x))dx$.
2. The expected welfare from time t on is given by $W_t = \sum_{i=1}^{k_t} q_{(i)} y_i(t)$, where $q_{(i)}$ is the i'th highest element of Π_t.

Albright: Illustration.

Example

There are three objects; $\lambda = 1$; the distribution of agents' types is $F(x) = 1 - e^{-x}$. The following figure depicts the solution for $T = 5$:

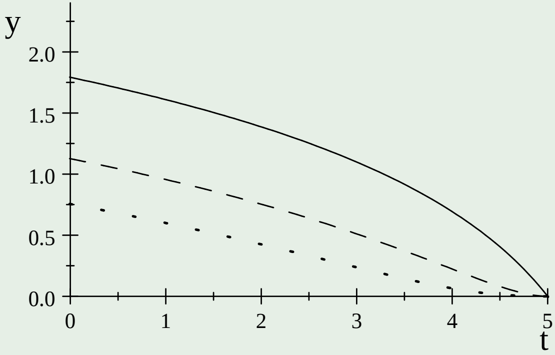

Second Order Stochastic Dominance I

Theorem

Consider two distributions of agents' types F and G such that $\mu_F = \mu_G = \mu$ and such that F second-order stochastically dominates G (in particular F has a lower variance than G). Then it holds that:

1. *For any time t, and for any n, the vector $\{y_i^F(t)\}_{i=1}^n$ of optimal cutoffs under F is weakly sub-majorized by the vector $\{y_i^G(t)\}_{i=1}^n$ of optimal cutoffs under G.*
2. *For any time t and for any set of available objects at t, $\Pi_t \neq \emptyset$, the expected welfare in the efficient dynamic allocation under F is lower than that under G.*

SSD: Illustration

Example

Let $F(x) = x$ on $[0, 1]$ so that F is IFR and thus F SSD $G(x) = 1 - e^{-2x}$. There is one object; $\lambda = 1$; $T = 5$. $y_1^F(t) = 1 - \frac{2}{7-t}$ is the dashed line, $y_1^G(t) = \frac{1}{2}\ln[6 - t]$ is the solid line.

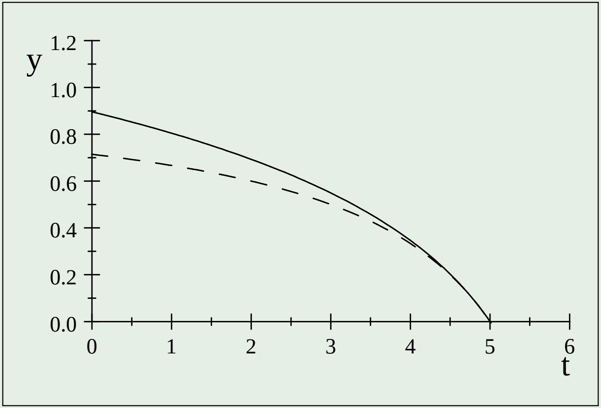

Second Order Stochastic Dominance (Static Counterpart)

Theorem (De La Cal & Caracamo, J. Appl. Prob. 06)

Consider two distributions of agents' types F and G. The following two assertions are equivalent:

1. *F second-order stochastically dominates G.*
2. *For any n, the $n-$tuple $\{\mu^F_{(i),n}\}^n_{i=1}$ of mean order statistics under F is majorized by the $n-$tuple $\{\mu^G_{(i),n}\}^n_{i=1}$ of mean order statistics under G.*

Corollary

Consider two distributions of agents' types F and G such that F second-order stochastically dominates G. Then, for any set of available objects, expected welfare in the efficient static allocation under F is lower than that under G.

The Loss from Sequentiality I

- Assume that at time t there are n objects left.
- Consider scenario where the allocation to all subsequently arriving agents can be made at the deadline T.
- Expected welfare at t is given by $\sum_{i=1}^{n} q_{(i)} z_i(t)$, where $z_i(t)$ represents the expected type of an agent who arrives after t, and who get assigned to the object with the $i-th$ highest quality.
- Let $\Pr_l(t) = e^{-\lambda(T-t)} \frac{\lambda^l (T-t)^l}{l!}$ be the probability that there will be l arrivals, $l \geq 1$, after time t. Then

$$z_i(t) = \sum_{l=i}^{\infty} \Pr_l(t) \mu_{(l-i+1),l} \; = e^{-\lambda(T-t)} \sum_{l=i}^{\infty} \frac{\lambda^l (T-t)^l \mu_{(l-i+1),l}}{l!}, \; \forall i.$$

where let $\mu_{(i),l}$ is the expectation of the $i-th$ highest order statistic out of l copies of X.

The Loss from Sequentiality II

- Intuitively, $\sum_{i=1}^{n} q_{(i)}[y_i(t) - z_i(t)]$ measures the welfare loss due to the sequentiality constraint.

Theorem

1. *For any period* t*, and for any* n*, the vector* $\{y_i(t)\}_{i=1}^{n}$ *of optimal cutoffs in the dynamically efficient allocation of* n *objects is weakly sub-majorized by the vector* $\{z_i(t)\}_{i=1}^{n}$*.*
2. *Moreover,* $\lim_{n\to\infty} \sum_{i=1}^{n} y_i(t) = \lim_{n\to\infty} \sum_{i=1}^{n} z_i(t) = \lambda(T - t)\mu$*, where* μ *is the mean of the distribution of agents' types.*

Infinite Horizon & Discounting

- Infinite horizon $(T = \infty)$, discount rate $r(t) = e^{-\alpha t}$.
- Arrival process is a *renewal* with general inter-arrival distribution B

Theorem (Albright MS 1974)

The dynamically efficient cutoff curves are stationary (i.e., independent of time) $y_n \leq y_{n-1} ... \leq y_1$. These constants do not depend on the q's, and are given by the implicit recursion:

$$(y_k + y_{k-1} + ..y_1) = \frac{\mathcal{L}_B(\alpha)}{1 - \mathcal{L}_B(\alpha)} \int_{y_k}^{\infty} (1 - F(x))dx \ , \ 1 \leq k \leq n$$

where $\mathcal{L}_B$ is the Laplace-transform of the inter-arrival distribution.

- Recall that for random variable X with density f, the *Laplace Transform* is given by $\mathcal{L}(s) = E[e^{-sX}]$

Inter-Arrival Distributions I

Definition

Let β, γ be two non-negative random variables governed by distributions B, C, respectively. Then B dominates C in the *Laplace transform order* if

$$E[e^{-s\beta}] \geq E[e^{-s\gamma}] \text{ for all } s > 0$$

Theorem

Consider two inter-arrival distributions B and C such that $B \geq_{Lt} C$.

1. *For any fixed distribution of agents' characteristics F, and for any n, the $n-$tuple $\{y_i^B\}_{i=1}^n$ of optimal cutoffs under B is weakly sub-majorized by the $n-$tuple vector $\{y_i^C\}_{i=1}^n$ of optimal cutoffs under C.*
2. *In particular, for any t, and for any $\Pi_t \neq \emptyset$, the expected welfare from t on in the efficient dynamic allocation under B is lower than that under C.*

Inter-Arrival Distributions II

Definition

A non-negative random variable W is called *NBU (NWU)* if, for every $y > 0$, W is stochastically larger (smaller) than the conditional random variable $(W - y / W \geq y)$.

Example

Assume that there is one object, and consider a situation with an NBU (NWU) distribution of abilities with mean μ, and another NBU (NWU) distribution of inter-arrival times with mean ω. Then, the expected welfare under the efficient dynamic policy is lower (higher) than

$$\mu \, \mathrm{LambertW}(\frac{1}{\omega\alpha})$$

where the increasing function $\mathrm{LambertW}(x)$ is implicitly defined by:

$$\mathrm{LambertW}(x) e^{\mathrm{LambertW}(x)} = x$$

Mechanisms and Incomplete Info. I

- Assume types are private information. If an item with quality $q_i \geq 0$ is assigned to an agent with type x_j for price p_i, this agent has utility $q_i x_j - p_i$.
- W.l.o.g. restrict attention to deterministic, Markovian and direct mechanisms where every agent, upon arrival, reports his characteristic x_i and where, at any point in time t, the mechanism specifies:

1. a non-random allocation rule (which object is allocated, if any) that only depends on t, on the declared type of the arriving agent, and on the inventory of items available at t.
2. a payment to be made by the arriving agent which depends on t, on the declared type of the agent, and on the inventory of items available at t.

Mechanisms II

Theorem

1) *An non-randomized, Markovian policy* $\{Q_t(x,\Pi_t)\}_t$ *is implementable iff at each t it partitions the set of agents' types into* k_t+1 *disjoint intervals such that all types in a given interval obtain the same quality, and such that higher types obtain a higher quality.*
2) *The associated payment scheme is given by*

$$P_t(x,\Pi_t)=\sum_{i=j}^{k_t}(q_{(i:\Pi_t)}-q_{(i+1:\Pi_t)})y_{i,\Pi_t}(t)\ \text{ if } x\in[y_{j,\Pi_t}(t),y_{j-1,\Pi_t}(t))$$

and by zero otherwise.

Mechanisms III

Corollary

The WM policy is implementable. Payments are the dynamic analogue of the Vickrey-Clarke-Groves mechanism (see also Bergemann and Valimäki, EC 10)

- Payments have intuitive interpretation: look at static case with k_t objects and $k_t + 1$ agents where, in addition to agent with type x, there are k_t "dummies" with types $y_{1,\Pi_t}(t)$, $y_{2,\Pi_t}(t)$, ..., $y_{k_t,\Pi_t}(t)$.
- Payment for object with the j-highest quality, $\sum_{i=j}^{k_t}(q_{(i:\Pi_t)} - q_{(i+1:\Pi_t)})y_{i,\Pi_t}(t)$ represents the externality imposed by agent with type x on dummies

RM with Heterogenous Objects

Theorem

Assume that the virtual value $x - \frac{1-F(x)}{f(x)}$ *is increasing. Then the RM allocation is given by n cut-off functions that do not depend on the available qualities. These functions satisfy:*

$$y_i(t) - \frac{1 - F(y_i(t))}{f(y_i(t))} + \lambda \int_t^T \frac{[1 - F(y_{i-1}(s))]^2}{f(y_{i-1}(s))} ds = \lambda \int_t^T \frac{[1 - F(y_i(s))]^2}{f(y_i(s))}$$

or, equivalently

$$y_i(t) - \frac{1 - F(y_i(t))}{f(y_i(t))} + R(\mathbf{1}_{i-1}, t) = R(\mathbf{1}_i, t)$$

where $R(\mathbf{1}_j, t)$ *is the expected revenue at time t from the optimal policy if j identical objects with* $q = 1$ *are still available at that time.*

- One object with quality q is available at time t ; agent arriving at t gets object if $x_i \geq y_1(t)$, and no object otherwise. Expected revenue given by

$$q \int_t^T y_1(s) h_1(s) ds$$

where $h_1(s)$ is the density of waiting time till the first arrival of an agent with a value above $y_1(s)$.

- By the *Coloring Theorem*, this density equals the density of the first arrival in a non-homogenous Poisson process with rate $\lambda(1 - F(y_1(s))$, given by

$$h_1(s) = \lambda(1 - F(y_1(s))e^{-\lambda \int_t^s [1-F(y(z))]dz} \text{ for } t \leq s \leq T$$

- Perform calculus of variation exercise with respect to $y_1(t)$
- Use backward induction.

Clearance Sales

- Percentage Markdown: Difference between prices of the same product at $t = 0$ and $t = T$, divided by the price at $t = 0$.
- Pashigian and Bowen (*QJE 91*) empirically find that :

"More expensive apparel items within each product line are frequently sold at a higher average percentage markdown"

Theorem

Assume an RM seller and consider the scenario where at time $t = 0$ there are $n_1 > 0$ items of quality q and $n_2 > 0$ items of quality $s < q$, while at time $t = T$ there are $l_1 > 0$ items of quality q , and $l_2 > 0$ items of quality s left unsold. Then the percentage markdown is always higher for the higher quality.

Optimal Inventory Design

Theorem

Let $y = \{y_i(t)\}_{i=1}^{n}$ *denote the allocation underlying the RM policy with* n *objects, and assume that the cost of producing qualities* $(q_1, q_2, ..q_n)$ *is given by* $C(q_1, q_2, ..q_n) = \sum_{i=1}^{n} \phi(q_i)$ *where* $\phi : R \longrightarrow R$ *is strictly increasing, convex and satisfies* $\phi(0) = 0$. *Then*

1) The optimal number of objects n^* *is characterized by*

$$\phi^{'}(0) \in (y_{n^*+1}(0) - \frac{1 - F(y_{n^*+1}(0))}{f(y_{n^*+1}(0))},\ y_{n^*}(0) - \frac{1 - F(y_{n^*}(0))}{f(y_{n^*}(0))}]$$

2) The optimal qualities q_i^* *are given by:*

$$\phi^{'}(q_i^*) = y_i(0) - \frac{1 - F(y_i(0))}{f(y_i(0))},\ i = 1, ..., n^*$$

Multi-Unit Demand: Dynamic Knapsack

- An RM seller has capacity $C \in \mathbb{R}_+$ that perishes after T periods.
- In each period, impatient agent arrives with quantity request w, and per-unit value v. Type (w, v) is private information to the arriving agent.
- Type (w, v)'s utility is given by $wv - p$ if at price p he is allocated a capacity $w' \geq w$ and by $-p$ if he is assigned an insufficient capacity $w' < w$.
- Demands are I.I.D across periods, governed by c.d.f. $F(w, v)$ with density $f(w, v) > 0$ on $[0, \infty)^2$. For all w, the conditional virtual value $v - \frac{1-F(v|w)}{f(v|w)}$ is an unbounded, strictly monotone function of v.
- Complete information optimization model due to Kleywegt & Papastavrou, *OR 01;* they do not consider payments.

Dynamic Knapsack: Implementable Policies

- A deterministic, Markovian allocation rule for time t with remaining capacity c has the form $\alpha_t^c : [0, +\infty)^2 \rightarrow \{1, 0\}$ where 1 (0) means that the reported capacity demand w is satisfied (not satisfied).

Theorem

A policy $\{\alpha_t^c\}_{t,c}$ is implementable iff for every t and every c it satisfies:
1) $\forall\, (w, v), v' \geq v,\ \alpha_t^c(w, v) = 1 \Rightarrow \alpha_t^c(w, v') = 1.$
2) The function $wp_t^c(w)$ is non-decreasing in w, where
$p_t^c(w) = \inf\{v \ / \ \alpha_t^c(w, v) = 1\}.$
The maximal, individually rational payment function that implements $\{\alpha_t^c\}_{t,c}$ is given by

$$q_t^c(w, v) = \begin{cases} wp_t^c(w) & \text{if } \alpha_t^c(w, v) = 1 \\ 0 & \text{if } \alpha_t^c(w, v) = 0 \end{cases}$$

Illustration: Implementable Policies

Example

Assume $T = 1$. Weight w is realized according to an exponential distribution with parameter λ. Per-unit value is sampled from the following distribution

$$F(v|w) = \begin{cases} 1 - e^{-\overline{\lambda} v} & \text{if} \quad w > w^* \\ 1 - e^{-\underline{\lambda} v} & \text{if} \quad w \leq w^* \end{cases}$$

where $\overline{\lambda} > \underline{\lambda}$ and $w^* \in (0, c)$. For observable weight requests, seller charges $\frac{1}{\underline{\lambda}}$ $(\frac{1}{\overline{\lambda}})$ per unit if the $w \leq (\geq)$ w^*. This implies that

$$wp_t^c(w) = \begin{cases} \frac{w}{\overline{\lambda}} & \text{if} \quad w > w^* \\ \frac{w}{\underline{\lambda}} & \text{if} \quad w \leq w^* \end{cases} .$$

and therefore, $wp_t^c(w)$ is not monotone.

Dynamic Knapsack: RM I

Theorem

Assume that:

1. *For any w, the hazard rate $\frac{f(v|w)}{1-F(v|w)}$ is non-decreasing in v.*
2. *For any $w' \geq w$, and for any v, $\frac{f(v|w)}{1-F(v|w)} \geq \frac{f(v|w')}{1-F(v|w')}$.*
   For each c, t, w let $p_t^c(w)$ denote the recursive solution to the system

$$w\left(p_t^c(w) - \frac{1-F(p_t^c(w)|w)}{f(p_t^c(w)|w)}\right) = R^*(c, T-t) - R^*(c-w, T-t).$$

where R^* denotes the optimal revenue with $R^*(c,0) = 0$ for all c. Then the underlying allocation where $\alpha_t^c(w,v) = 1$ iff if $v \geq p_t^c(w)$ is implementable. In particular, the above system determines the RM policy.

Dynamic Knapsack: RM II

- Optimal revenue may not be concave in capacity (see Kleywegt & Papastavrou, *OR 01* for non-concavity in WM) - this is connected to implementation problems.
- Under a concavity conditions on the distribution of types, revenue is concave and above system also determines RM policy.
- RM policy requires price adjustments for every c, t, w - complicated dynamics. We construct a static nonlinear price schedule (it uses correlations between w and v !) that is asymptotically optimal if $\min(C, T)$ is large enough.

Learning about Values: Illustration 1

- one object
- two agents arrive sequentially, one per period
- each agent can only be served upon arrival
- after an item is assigned, it cannot be reallocated
- valuations x_i are private
- valuations distributed independently and uniformly on $[0, 2]$

Illustration 1: continued

The dynamically efficient allocation is given by

- first agent gets the object if $x_1 \geq 1$
- second agent gets the object if $x_1 < 1$
- Implemntation with payment scheme:
- for the first player: $P_1\left(x_1\right) = \begin{cases} 1 & \text{if} \quad x_1 \in [1,2] \\ 0 & \text{if} \quad x_1 \notin [1,2] \end{cases}$
- for the second player: $P_2\left(x_2\right) = 0$

Learning about Values: Illustration 2

- one object
- two agents arrive sequentially, one per period
- each agent can only be served upon arrival
- the designer does not know the distribution of values
- with probability 0.5 the distribution of both agents' types is uniform on $[0, 1]$
- with probability 0.5 the distribution is uniform $[1, 2]$

Value of keeping vs. value of allocating object

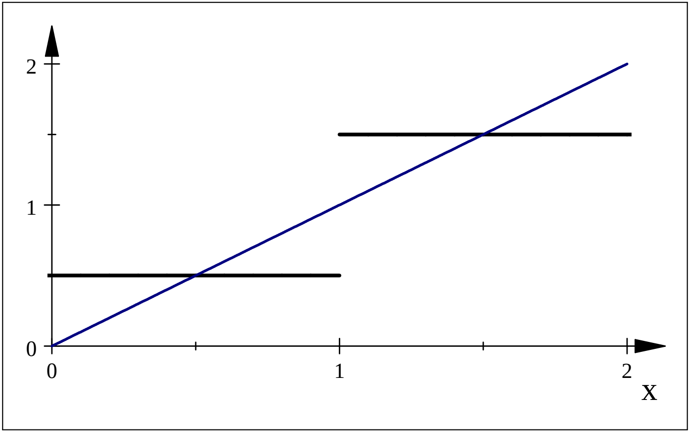

Illustration 2: continued

- The dynamically efficient allocation is
- first agent gets the object if $x_1 \in [0.5, 1] \cup [1.5, 2]$
- second agent gets the object if $x_1 \notin [0.5, 1] \cup [1.5, 2]$
- The payment scheme for the first agent (P, p) should satisfy:

$$\begin{aligned} x_1 + P &\geq p \quad \text{for any} \quad x_1 \in [0.5, 1] \\ x_1 + P &\leq p \quad \text{for any} \quad x_1 \in [1, 1.5] \end{aligned}$$

- But this is IMPOSSIBLE!

Complete Information & Learning

Theorem (Albright, 1977)

Consider the arrival of an agent with type x in period $k \geq 1$, and denote by χ_k the vector of past reported signals. There exist $k+1$ functions

$$0 = a_{0,k}\left(\chi_k, x)\right) \leq a_{1,k}\left(\chi_k, x\right) \leq a_{2,k}\left(\chi_k, x\right) \leq \ldots \leq a_{k,k}\chi_k, x_k\right) = \infty$$

such that the dynamically efficient policy assigns the item with the i'th smallest type if $x \in \left(a_{i-1,k}\left(\chi_k, x\right), a_{i,k}\chi_k, x)\right]$.

- *Each $a_{i,k+1}(\chi_{k+1}, x_{k+1})$ equals the expected value of the agent's type to which the item with $i-th$ smallest type is assigned in a problem with k periods before the period k signal is observed. These functions do not depend on the q's.*

Incomplete Information & Learning-I

Question: When do cutoffs exist which

1. are independent of the signal of the current agent and
2. replicate the efficient allocation ?

The key element is the set

$$A_{i,k}(\chi_k) = \{x : a_{i-1,k}(\chi_k, x)) \leq x < a_{i,k}(\chi_k, x))\}.$$

B 4 - Incomplete Information & Learning-II

Theorem (Necessary and sufficient condition)

The efficient allocation is implementable if and only if for any period k, any object q_i and any history H_k the set $A_{i,k}(\chi_k)$ is convex.

Theorem (Sufficient condition)

Assume that for any k, χ_k, $i \in \{0, .., k\}$, the cutoff $a_{i,k}(\chi_k, x_k)$ is a Lipschitz function of x_k with constant 1. Then, the efficient dynamic policy is implementable.

- Insight taken from the theory of mechanism design with interdependent values (Jehiel & Moldovanu, *EC 01*)

Sufficient Conditions I

Theorem

Assume that for any k, and for any pair of ordered lists of reports $\chi_k \geq \chi'_k$ that differ only in one coordinate, the following conditions hold:

1. $\widetilde{F}_k(x|\chi_k) \succsim_{FOSD} \widetilde{F}_k(x|\chi'_k)$
2. $E(x|\chi_k) - E(x|\chi'_k) \leq \frac{\Delta}{k-1}$ *where* Δ *is size of the difference between* χ_k *and* χ'_k

   Then, the efficient dynamic policy is implementable.

Sufficient Conditions II

Theorem

Assume that the conditional distribution function $\widetilde{F}_k\left(x|x_n, \cdots, x_{k+1}\right)$ *and density* $\widetilde{f}_k\left(x|x_n, ..., x_{k+1}\right)$ *are continuously differentiable with respect to* x_{k+i} *for all* x *and for all* $n-k \geq i \geq 1$*. If*

$$0 \geq \frac{\partial}{\partial x_{k+i}} \widetilde{F}_k\left(x|\chi_k\right) \geq -\frac{1}{n-k} \widetilde{f}_k\left(x|\chi_k\right) \text{ for all } x \text{ and all } n-k \geq i \geq 1$$

then the efficient dynamic policy is implementable.

Illustration

Example

Consider $\widetilde{x} \sim N(\mu, 1)$ with unknown mean μ, and prior beliefs $\widetilde{\mu} \sim N(\mu_0, \frac{1}{\tau})$ with $\tau > 0$.

- After observing $x_n, ..x_{k+1}$, posterior is $\widetilde{\mu} \sim N(\overline{\mu}, \frac{1}{\tau+n-k})$ with $\overline{\mu} = \frac{\tau\mu_0 + \sum x_i}{\tau+(n-k)}$, and

$$
\begin{aligned}
\widetilde{F}_k(x|x_n, ..., x_{k+1}) &= N(\overline{\mu}, 1 + \frac{1}{\tau + n - k}); \\
\widetilde{F}_k(x + \frac{z}{\tau + (n-k)}|x_n, .., x_i + z, .., x_{k+1}) &= \widetilde{F}_k(x|x_n, .., x_i, .., x_{k+1})
\end{aligned}
$$

- SD holds, and

$$
\begin{aligned}
\frac{\partial \widetilde{F}_k\left(x|x_n, .., x_{k+1}\right)}{\partial x_{k+i}} &= -\frac{1}{\tau + n - k}\widetilde{f}_k\left(x|x_n, .., x_{k+1}\right) \\
&\geq -\frac{1}{n-k}\widetilde{f}_k\left(x|x_n, .., x_{k+1}\right)
\end{aligned}
$$

Remarks

Remark

1. *The right-hand side condition says that the function* $\widetilde{F}_k\left(x+\frac{z}{n-k}\middle|x_{k+1},..,x_{k+i}+z,x_{k+i+1},..,x_n\right)$ *is non-decreasing in* z*. In other words, after having made* $n-k$ *observations, a small shift to the right - which moves the value of the distribution upwards - is enough to compensate the downward effect on the distribution's value caused by an* $(n-k)$ *times larger upward shift in one of the past observations (recall that, by stochastic dominance, higher observations move the entire distribution downwards).*
2. *The condition guarantees that* $\forall i,k,n,\ \frac{\partial a_{i,k}(\chi_k,x_k)}{\partial x_k}\leq\frac{1}{n-k}$ *although* $\frac{\partial a_{i,k}(\chi_k,x_k)}{\partial x_k}\leq 1$ *seems sufficient for the implementation of the efficient allocation. Nevertheless, the long-term effect of each non-terminal observation makes it impossible to obtain tighter conditions that apply generally.*

Search for the Lowest Price (Rothschild1973):

- Consumer obtains a sequence of prices, and must decide when to stop the (costly) search for a lower price.
- Beliefs about distribution of prices updated (in a Bayesian way) after each observation.
- Without learning, optimal stopping rule is characterized by a reservation price R : stops search at any price less than equal to R, and continue search at any price higher than R.
- If all customers follow such a rule $\Longrightarrow$ well-behaved demand function where sales are a non-increasing function of price
- When does the optimal stopping rule with learning have the *reservation price property* ? I.e.,when does a price $R(s)$ exist for each state s such that prices above are rejected, and prices below are accepted ?

Answers to Rothschild's Question

- Stochastic Dominance $\dot{+}$:
- Rosenfield and Shapiro (1981): For all x, k, χ_k and all $n-k \geq i \geq 1$

$$\int_x^\infty \frac{\partial}{\partial x_{k+i}} \widetilde{F}_k\left(y|\chi_k\right) dy \geq -\frac{1}{n-k}(1-\widetilde{F}_k(x|\chi_k)$$

- Seierstad (1992) : For all x, k and χ_k

$$\sum_{i=1}^{n-k} \frac{\partial}{\partial x_{k+i}} \widetilde{F}_k\left(x|\chi_k\right) \geq -\widetilde{f}_k\left(x|\chi_k\right)$$

- Us (multiple objects !): For all x , χ_k ,and all $n-k \geq i \geq 1$

$$\frac{\partial}{\partial x_{k+i}} \widetilde{F}_k\left(x|\chi_k\right) \geq -\frac{1}{n-k} \widetilde{f}_k\left(x|\chi_k\right)$$

Non-Bayesian Learning I

Before stage n the designer's prior belief about the distribution of the first type x_n is given by G. Conditional on observing $x_n, x_{n-1} \ldots, x_{k+1}$ at stages $n, n-1, \ldots k+1$, the belief about the distribution of the next type $x = x_k$ is given by:

$$\widetilde{F}_k(x|x_n, \ldots, x_{k+1}) = (1 - \beta_k^n) G(x) + \beta_k^n \frac{1}{n-k} \sum_{i=k+1}^{n} \mathbf{1}_{[x_i, \infty)}(x), \ k = 1, 2, \ldots n -$$

where $0 < \beta_k^n < 1$ and where $\mathbf{1}_{[z,\infty)}(x)$ denotes the indicator function of the set $[z, \infty)$.

Theorem

Under "empirical" learning , the efficient dynamic policy can always be implemented under incomplete information.

Non-Bayesian Learning II

Prior to any observation, the designer estimates the unknown distribution to be uniform. Suppose that m observations were made, and order them in increasing order $\left\{x_{(1)}, \ldots, x_{(m)}\right\}$, and let $x_{(0)} = 0$ and $x_{(m+1)} = 1$. The belief about the distribution of the next type x is given by:

$$f_k\left(x | x_n, \ldots, x_{n-m+1}\right) = \sum_{i=1}^{m+1} \frac{\mathbf{1}_{[x_{(i-1)}, x_{(i)})}(x)}{(m+1)\left(x_{(i)} - x_{(i-1)}\right)}.$$

Theorem

Under "maximum entropy" learning, the efficient dynamic policy can always be implemented under incomplete information.

Characterization of Second-Best Mechanism

Theorem

1. *At each period k, expected welfare (calculated before the arrival of the period k agent) is a Schur-convex, linear function of the available qualities at that period.*
2. *The incentive compatible, optimal mechanism (second best) is deterministic. That is, for every history at period k, H_k, and for every type x_k of the agent that arrives at that period, there exists a quality q that is allocated to that agent with probability 1 .*
3. *At each period, the optimal mechanism partitions the type set of the arriving agent into a collection of disjoint intervals such that all types in a given interval obtain the same quality with probability 1, and such that higher types obtain a higher quality.*

Outline of Proof

- Schur-convexity of expected welfare in a deterministic, incentive compatible mechanisms follows by induction.
- $\Rightarrow$ At each period k it is optimal to leave for the future the "most disperse" set (in the sense of majorization) of feasible qualities that is consistent with incentive compatibility.
- $\Rightarrow$ Period k's optimal allocation must be the "most concentrated" one that is consistent with monotonicity.
- $\Rightarrow$ Period k's allocation should be either deterministic, or should randomize among at most two neighboring qualities.
- Finally, randomization among two neighboring qualities cannot be optimal. This is related to an argument due to Riley & Zeckhauser who studied revenue maximization for 1 seller-1 buyer problem without monotone virtual values.

Patient Agents & Complete Info. (Bertsekas 05)

- Seller has 1 object.
- Time is discrete, deadline T.
- One new buyer arrives at each period t (arrivals observed by seller).
- Buyers have IID values for the object represented by random variable $\widetilde{v}$
- Buyers are long lived; common discount factor δ.

Patient Agents: Main Result

Theorem

The dynamically efficient policy awards the good in period t if and only if the set of present buyers contains a buyer whose value exceeds a cutoff x_t. If the good is allocated in period t , it is awarded to the agent with the highest value. The cutoffs x_t satisfy:

1. $x_T = 0$
2. $x_t = \delta E[\max(\widetilde{v}, x_t)]$, *for* $t < T$. *In particular the optimal cutoff is constant in periods* $t < T$.

- Intuition: At the optimum in period $t < T$, the seller should be indifferent between awarding the good to a buyer with value x_t, and waiting one more period and getting another draw. Note that x_t, $t > T$, also equals the optimal cutoff in the infinite horizon problem with **impatient** agents !

Patient Agents: Incomplete Information (Board & Skrzypacz 10)

- Assume now that values are private information.
- Optimal allocation can be implemented through a series of second-price auctions.
- Although optimal cutoff is constant for all $t < T$, reserve prices that implement these cutoffs decrease over time !
- The reserve price makes the cutoff-type indifferent between buying today and tomorrow; Prices thus decline both because discounting and because the risk of losing to a new bidder that comes tomorrow.
- Agents with value below cutoff but above reserve price refrain from buying: they wait for the "fire-sale" at the last period.

Strategic and Unobservable Arrivals

- Designer endowed with indivisible object;
- Stream of randomly arriving, **long-lived** (or patient) agents; arrivals described by a counting process $\{\mathcal{N}(t), t \geq 0\}$ in continuous time
- Agent's private information is two-dimensional: arrival time $t \geq 0$ and value $v \geq 0$ for the object. Arrival time and value are independent of each other.
- Values are I.I.D. random variables with common c.d.f. F, continuos density f.
- If an agent arrives at t, gets the object at $\tau \geq t$ and pays p at $\tau' \in [t, \tau]$, then her utility is given by $e^{-\delta\tau}v - e^{-\delta\tau'}p$ where $\delta \in (0, 1)$ is the discount factor.
- Designer's utility is given by $e^{-\delta\tau}v$ (WM) or $e^{-\delta\tau'}p$ (RM).

Renewals under Complete Information

Theorem (Zuckerman)

A non-negative random variable W is called NBU (new better than old) if, for every $y > 0$, W is stochastically larger than the conditional random variable $(W - y/W \geq y)$.

Assume that the arrival process is a renewal, and that the inter-arrival distribution G satisfies the NBU property. Let ϕ denote the Laplace Transform of G. Then, the optimal policy allocates the object to the first arrival whose value is above v^* where v^* is the unique solution to

$$v^* = \frac{\phi(\delta)}{1-\phi(\delta)} \int_{v^*}^{\infty} (v - v^*) dF(v) dv.$$

In particular, recall is never used by the optimal policy, and all allocations occur upon arrival.

Application to RM (Gallien 06)

- Private information is characterized by (v_i, t_i, d_i) where v_i is the value, t_i is the arrival time, and $d_i \geq t_i$ is the exit time. The designer's goal is to maximize revenue.
- Assume that arrivals are governed by a renewal process with a *NBU* distribution, and that the virtual valuation function $x - \frac{1-F(x)}{f(x)}$ is strictly increasing, and denote by H the implied distribution of virtual values. Then, for any d_i, $d_i \geq t_i$, $i = 1, 2, ..$, the revenue maximizing policy is to charge a constant price P where P is the unique solution to
$$p = \frac{\phi(\delta)}{1 - \phi(\delta)} \int_p^\infty (z - p) dH(z) dz$$

The solution coincides with the one found by Gallien for the case $d_i = \infty$ (perfect recall) and by GM for the case $d_i = t_i$ (no recall).

Learning from Arrivals: Example

- Distribution of values is $U\left[0,1\right]$.
- Inter-arrival time is known: 1) $U[1,2]$ or 2) $U\left[2,3\right]$.
- The optimal cutoffs for a WM designer are

$$x_{[i,j]}(\delta) = \frac{1}{\phi_{[i,j]}\left(\delta\right)}\left(1 - \sqrt{1 - \left(\phi_{[i,j]}\left(\delta\right)\right)^2}\right)$$

where

$$\phi_{[i,j]}\left(\delta\right) = \frac{e^{-i\delta} - e^{-j\delta}}{\delta}.$$

is the respective Laplace transform.

- Fix any δ and note that $x_{[1,2]} = x_{[1,2]}\left(\delta\right) > x_{[2,3]}\left(\delta\right) = x_{[2,3]}$.
- No recall used

Unknown Arrival Process / Complete Information

- Designer does not know the distribution of inter-arrival times: believes that it is either $U[1,2]$ or $U[2,3]$, with equal probabilities.
- Under complete information optimal policy is:

1. For $T \in (1,2]$ the cutoff is $x_{[1,2]}$
2. For $T \in (2,3]$ the cutoff is $x_{[2,3]}$ if there were no arrivals before time 2, otherwise the cutoff is $x_{[1,2]}$
3. For $T > 3$, the cutoff is $x_{[1,2]}$ if the first arrival happened during time interval $(1,2]$, whereas the cutoff is $x_{[2,3]}$ if the first arrival happened during $(2,3]$.

Unknown Arrival Process / Incomplete Information

- The allocative externality payment, which needs to be paid for the object by an agent arriving at $t \geq t_1$ is

$$P(t_1) = \begin{cases} x_{[1,2]} & \text{if } t_1 \in [1,2] \\ x_{[2,3]} & \text{if } t_1 \in (2,3] \end{cases} .$$

- Consider type (t, v) with $t \in (1,2)$ and $v \in (x_{[2,3]}, x_{[1,2]})$. Truthful reporting yields zero utility since object not allocated to him. But, a report (t', v) where $t' = t + 1 \in (2,3)$ yields utility $e^{-\delta t'} \left(v - x_{[2,3]} \right) > 0$.
- Truthful reporting under standard Clarke-Groves-Vickrey mechanism is not optimal, and WM dynamic allocation cannot be implemented !
- Problem: Informational externality on designer induced by early arrivals.

Subsidizing Early Arrivals I

- Subsidy to be paid to agent that arrives at t, independently of whether he gets the object:

$$S(t) = \begin{cases} x_{[1,2]} - x_{[2,3]} > 0 & \text{if } t \in [1,2] \\ 0 & \text{if } t \geq 2 \end{cases} .$$

- General intuition: Let $U(t, t', v)$ be the utility of an agent with type (t, v) who arrives at t' and optimally reports v' (given (t', t, v)) in the version of the model where arrivals are observable and where the designer uses the dynamically efficient mechanism.
- Choose subsidy scheme $S(t')$ such that the function $U(t, t', v) + e^{-\delta t'} S(t')$ decreases in reported arrival time t' for any t and v.

Subsidizing Early Arrivals II

Theorem

Assume that there exists $M \geq 0$ such that for any $t \leq t' \leq t''$ in $[0, T]$ and for any v it holds that

$$U\left(t, t'', v\right) - U(t, t', v) \leq M(t'' - t')$$

Then the subsidy $S(t') = e^{\delta t'} M(T - t')$ together expected allocative externality payment implements the efficient policy for any history which stops before T.

- Let $L(t') = \sup_{v, t \leq t'} \frac{\partial U(t, t', v)}{\partial t'}$ and assume that there exists $c(t') \geq L(t')$ such that $C = \int_0^\infty c(z)dz < \infty$. The subsidy function

$$S(t') = e^{t'\delta}\left[C - \int_0^{t'} L(z)dz\right]$$

works for $T = \infty$.

Winner-Pay Mechanisms I

Definition

A mechanism is called *winner-pay mechanism* if the transfers to all agents that do not get the object are zero.

Theorem

Assume that arrivals are unobservable, and that the arrival process is a known renewal with inter-arrival distribution G. Then a subsidy is not needed for efficient implementation.

Proof.

(Sketch): Show that $U(t, t', v)$ decreases in t' for any v and for any $t \leq t'$.At allocation, elapsed time from last arrival must be the same, independently of reported arrival time. Price charged depends on elapsed time since last arrival, and on second highest value reported so far. Later arrival postpones allocation $\longrightarrow$ more arrivals $\longrightarrow$increases second highest value $\longrightarrow$ increases price. □

Winner-Pay Mechanisms II

Theorem

Assume that arrivals are governed by a non-homogenous Poisson/pure-birth process with arrival rate $\lambda_i(t)$ where λ is non-decreasing in both elapsed time t and the number of arrivals i up to t. Suppose that there exists a non-decreasing function $\beta(t)$ such that $\lambda_i(t) \leq \beta(t)$ for all i, t, and such that $\int_0^\infty \beta(t)e^{-\delta t}dt < \infty$. Then:

1. *An optimal stopping rule exists for the setting with observable arrivals.*
2. *Charging for the object a payment equal to the optimal cutoff under observable arrivals, implements the efficient allocation also if arrivals are non-observable. In other words, a subsidy is not needed for efficient implementation.*

Calculated Example

Example

Assume that the distribution of values is $F(v) = 1 - e^{-v}$, and consider a non-homogenous Poisson process with rate $\lambda(t) = \delta(t+2)\ln(t+2) - 1$. λ is positive and increasing in t, and

$$\int_0^\infty [\delta(t+2)\ln(t+2) - 1]e^{-\delta t} dt < \infty$$

Optimal cutoff $y(t)$ in the optimal stopping problem where arriving agents are short lived is given by:

$$y' - \delta y = -[\delta(t+2)\ln(t+2) - 1]e^{-y}$$

The solution $y(t) = \ln(t+2)$ increases in t and satisfies all optimality conditions for $\delta \geq \frac{1}{2\ln 2} = 0.72$.Charging $P(t) = \ln(t+2)$ implements the efficient dynamic allocation also in the problem with recall and unobservable arrivals if $\delta \geq \frac{1}{2\ln 2}$.

Impossibility Result: Pessimism Induced by Later Arrivals

Theorem

Consider an arrival process where the arrival rate at time t given that there were k arrivals before t is given by a differentiable function $\lambda_k(t_1, t_2, ..t_k, t) > 0$ where $t_1 \leq t_2.. \leq t_k \leq t$ are the k arrival times up to t. Assume that λ_k is strictly decreasing in $t_1, t_2, ..t_k$, k and t. Then, there is no winner-pay mechanism that implements the dynamically efficient allocation.

Long Lived Agents: Additional Topics

- Strategic deadlines, finite horizon (Pai & Vohra 08, Mierendorff 09)
- Queueing (Dolan, *Bell 78*, Hassin & Haviv 02, Kittsteiner & Moldovanu, *MS* 05)
- Changing information (types) over time (see Courty & Li 00, Battaglini 05, Pavan, Segal & Toikka 09)
- Short lived objects/long lived agents (Said 09)

Conclusion

- Modern theory of revenue management.
- Fruitful blend of

1. Dynamic models from the OR, MS, CS , Econ (search) literatures with focus on "grand, centralized optimization" and/or "ad-hoc", intuitive mechanisms.
2. Mechanism design literature with focus on information/incentives in static settings.

- Much remains to be done, e.g.

1. competition and dynamic pricing.
2. Dynamic pricing without commitment.