Individualized Pricing for a Cloud Provider Hosting Interactive Applications

Online Supplement

A Proof of Propositions

Proposition 2

Proof. The objective is to decompose w_j into μ_{rjt}s to minimize $\overline{V_t}(\mathbf{X})$ given $v_{r,t+1}$.

$$\begin{aligned}
\overline{V_t}(\mathbf{X}) &= \sum_{r=1}^{R} v_{r,t}(x_r) \\
&= \sum_{r=1}^{R} \mathbb{E}_j \left[\max\{\mu_{rjt} + v_{r,t+1}(x_r - \ell_r), v_{r,t+1}(x_r)\} \right] \\
&= \sum_{r=1}^{R} \mathbb{E}_j \left[v_{r,t+1}(x_r - \ell_r) \right] + \sum_{r=1}^{R} \mathbb{E}_j \left[\max\{\mu_{rjt}, v_{r,t+1}(x_r) - v_{r,t+1}(x_r - \ell_r)\} \right] \\
&= \sum_{r=1}^{R} \mathbb{E}_j \left[v_{r,t+1}(x_r - \ell_r) \right] + \mathbb{E}_j \left[\sum_{r=1}^{R} \max\{\mu_{rjt}, v_{r,t+1}(x_r) - v_{r,t+1}(x_r - \ell_r)\} \right]
\end{aligned}$$

The term $\sum_{r=1}^{R} \mathbb{E}_j \left[v_{r,t+1}(x_r - \ell_r) \right]$ is independent of the decomposition μ_{rjt}. We show that the term $\sum_{r=1}^{R} \max\{\mu_{rjt}, v_{r,t+1}(x_r) - v_{r,t+1}(x_r - \ell_r)\}$ is minimized when $\mu_{rjt} \propto v_{r,t+1}(x_r) - v_{r,t+1}(x_r - \ell_r)$.

Assume that for some j, t, r and r' (wlog let $r = 1$ and $r' = 2$), the contributions μ_{1jt} and μ_{2jt} are chosen such that

$$\frac{\mu_{1jt}}{v_{1,t+1}(x_1) - v_{1,t+1}(x_1 - \ell_1)} > \frac{\mu_{2jt}}{v_{2,t+1}(x_2) - v_{2,t+1}(x_2 - \ell_2)}.$$

We show that by subtracting d from μ_{1jt} and adding it to μ_{2jt} such that

$$\begin{aligned}
\frac{\mu_{1jt} - d}{v_{1,t+1}(x_1) - v_{1,t+1}(x_1 - \ell_1)} &= \frac{\mu_{2jt} + d}{v_{2,t+1}(x_2) - v_{2,t+1}(x_2 - \ell_2)} \\
&= \frac{\mu_{1jt} + \mu_{2jt}}{v_{1,t+1}(x_1) - v_{1,t+1}(x_1 - \ell_1) + v_{2,t+1}(x_2) - v_{2,t+1}(x_2 - \ell_2)},
\end{aligned} \tag{25}$$

the term $\max\{\mu_{1jt}, v_{1,t+1}(x_1) - v_{1,t+1}(x_1 - \ell_1)\} + \max\{\mu_{2jt}, v_{2,t+1}(x_2) - v_{2,t+1}(x_2 - \ell_2)\}$ does not increase. Formally:

$$\begin{aligned}
&\max\{\mu_{1jt} - d, v_{1,t+1}(x_1) - v_{1,t+1}(x_1 - \ell_1)\} + \max\{\mu_{2jt} + d, v_{2,t+1}(x_2) - v_{2,t+1}(x_2 - \ell_2)\} \\
\leq &\max\{\mu_{1jt}, v_{1,t+1}(x_1) - v_{1,t+1}(x_1 - \ell_1)\} + \max\{\mu_{2jt}, v_{2,t+1}(x_2) - v_{2,t+1}(x_2 - \ell_2)\}.
\end{aligned} \tag{26}$$

To show this inequality, we consider two possible cases:

Case 1: $\mu_{1jt} + \mu_{2jt} < v_{1,t+1}(x_1) - v_{1,t+1}(x_1 - \ell_1) + v_{2,t+1}(x_2) - v_{2,t+1}(x_2 - \ell_2)$. In this case the rightmost side of (25) is less than 1, therefore:

$$\begin{aligned}
&\max\{\mu_{1jt} - d, v_{1,t+1}(x_1) - v_{1,t+1}(x_1 - \ell_1)\} + \max\{\mu_{2jt} + d, v_{2,t+1}(x_2) - v_{2,t+1}(x_2 - \ell_2)\} \\
&= v_{1,t+1}(x_1) - v_{1,t+1}(x_1 - \ell_1) + v_{2,t+1}(x_2) - v_{2,t+1}(x_2 - \ell_2) \\
&\leq \max\{\mu_{1jt}, v_{1,t+1}(x_1) - v_{1,t+1}(x_1 - \ell_1)\} + \max\{\mu_{2jt}, v_{2,t+1}(x_2) - v_{2,t+1}(x_2 - \ell_2)\}.
\end{aligned}$$

Case 2: $\mu_{1jt} + \mu_{2jt} \geq v_{1,t+1}(x_1) - v_{1,t+1}(x_1 - \ell_1) + v_{2,t+1}(x_2) - v_{2,t+1}(x_2 - \ell_2)$. In this case, the rightmost side of (25) is ≥ 1, therefore:

$$\begin{aligned}
&\max\{\mu_{1jt} - d, v_{1,t+1}(x_1) - v_{1,t+1}(x_1 - \ell_1)\} + \max\{\mu_{2jt} + d, v_{2,t+1}(x_2) - v_{2,t+1}(x_2 - \ell_2)\} \\
&= \mu_{1jt} - d + \mu_{2jt} + d = \mu_{1jt} + \mu_{2jt} \\
&\leq \max\{\mu_{1jt}, v_{1,t+1}(x_1) - v_{1,t+1}(x_1 - \ell_1)\} + \max\{\mu_{2jt}, v_{2,t+1}(x_2) - v_{2,t+1}(x_2 - \ell_2)\}.
\end{aligned}$$

In both cases, the inequality (26) holds. Since this argument is true for any chosen r and r', we conclude that the optimal choice for the decomposed parameters μ_{rjt} is as described by the proposition:

$$\frac{\mu_{rjt}}{v_{r,t+1}(x_r) - v_{r,t+1}(x_r - \ell_{rjt})} = \frac{\mu_{r'jt}}{v_{r',t+1}(x_{r'}) - v_{r',t+1}(x_{r'} - \ell_{r'jt})} \quad \forall\; r, r' \in \{1, 2, ..., R\}$$

□

Proposition 3.

Proof. We begin with three lemmas.

Lemma 1. *Consider a T-period stochastic dynamic program $D1$ with state variable x_t, a period-specific stochastic variable z_t, action space $A(x_t, z_t)$, state transition function $\Gamma(x_t, a_t, z_t)$, revenue function $w(x_t, z_t, a_t)$, and value function $V_t(x_t) = \mathbb{E}[V_t(x_t, z_t)]$ given by the following Bellman equation:*

$$V_t(x_t, z_t) = \max_{a_t \in A(x_t, z_t)} w(x_t, z_t, a_t) + V_{t+1}(\Gamma(x_t, a_t, z_t)). \tag{27}$$

For any pair of intermediate periods t' and t'', the following inequality holds:

$$\begin{aligned}
\mathop{\mathbb{E}}_{\{z_t\}} \Big[\max_{\{a_t, x_t\}} \sum_{t=t'}^{t''-1} w(x_t, z_t, a_t) + V_{t''}(x_{t''}) \Big] \quad &\geq \quad V_{t'}(x_{t'}) \\
s.t.\;\; x_{t+1} = \Gamma(x_t, a_t, z_t)& \\
a_t \in A(x_t, z_t).&
\end{aligned} \tag{28}$$

Proof. Let $\hat{z}_t$ be the realizations of the stochastic variables z_t, and sequentially define $\hat{a}_t$ and $\hat{x}_t$, $t =$

$t', ..., t'' - 1$ by

$$\begin{aligned}
&\hat{a}_t = \arg \max_{\hat{a}_t \in A(\hat{x}_t, \hat{z}_t)} w(\hat{x}_t, \hat{z}_t, \hat{a}_t) + V_{t+1}(\Gamma(\hat{x}_t, \hat{a}_t, \hat{z}_t)) \\
&\hat{x}_{t+1} = \Gamma(\hat{x}_t, \hat{a}_t, \hat{z}_t),
\end{aligned} \tag{29}$$

then by definition

$$V_{t'}(x_{t'}) = \mathop{\mathbb{E}}_{\{z_t\}} \Big[\sum_{t=t'}^{t''-1} w(\hat{x}_t, \hat{z}_t, \hat{a}_t) + V_{t''}(\hat{x}_{t''}) \Big], \tag{30}$$

and the result follows by noting that $\hat{a}_t$ and $\hat{x}_t$ are feasible solutions to the maximization problem in (28). □

Lemma 2. *Consider a stochastic dynamic program D1 as defined in Lemma 1. Draw two instances of z_{t+1}: draw z^1 from the distribution of z_{t+1} (denoted by $\pi(z)$) and z^2 from a uniform distribution where all instances have equal probability. Define $\pi^1 \triangleq \frac{\pi(z^1)}{\pi(z^1)+\pi(z^2)}$ and $\pi^2 \triangleq \frac{\pi(z^2)}{\pi(z^1)+\pi(z^2)}$. The following inequality holds:*

$$\begin{aligned}
V_t(x_t, z_t) &= \max_{a_t \in A(x_t, z_t)} w(x_t, z_t, a_t) + \mathop{\mathbb{E}}_{z_{t+1}} [V_{t+1}(\Gamma(x_t, a_t, z_t), z_{t+1})] \\
&\leq \mathop{\mathbb{E}}_{z^1, z^2} \Big[\max_{a_t, a^1, a^2} w(x_t, z_t, a_t) + \pi^1 . \big(w(x_{t+1}, a^1, z^1) + V_{t+2}(\Gamma(x_{t+1}, a^1, z^1)) \big) \\
&\qquad\qquad + \pi^2 . \big(w(x_{t+1}, a^2, z^2) + V_{t+2}(\Gamma(x_{t+1}, a^2, z^2)) \big) \Big] \\
&\qquad\qquad s.t. \;\; a_t \in A(x_t, z_t), \;\; x_{t+1} = \Gamma(x_t, a_t, z_t), \;\; a^1, a^2 \in A(x_t, z_t).
\end{aligned} \tag{31}$$

Proof.

$$\begin{aligned}
V_t(x_t, z_t) &= \max_{a_t \in A(x_t, z_t)} w(x_t, z_t, a_t) + \mathop{\mathbb{E}}_{z_{t+1}} [V_{t+1}(\Gamma(x_t, a_t, z_t), z_{t+1})] \\
&= \max_{a_t \in A(x_t, z_t), a(z) \in A(x_{t+1}, z)} w(x_t, z_t, a_t) \\
&\quad + \sum_z \pi(z) \big(w(x_{t+1}, a(z), z) + V_{t+2}(\Gamma(x_{t+1}, a(z), z)) \big) \\
&= \max_{a_t, a(z) \text{ feasible}} w(x_t, z_t, a_t) \\
&\quad + \sum_{z^1, z^2} \Big(\frac{\pi(z^1) + \pi(z^2)}{J} \frac{\pi(z^1) \big(w(x_{t+1}, a(z^1), z^1) + V_{t+2}(\Gamma(x_{t+1}, a(z^1), z^1)) \big)}{\pi(z^1) + \pi(z^2)} \\
&\quad + \frac{\pi(z^1) + \pi(z^2)}{J} \frac{\pi(z^2) \big(w(x_{t+1}, a(z^2), z^2) + V_{t+2}(\Gamma(x_{t+1}, a(z^2), z^2)) \big)}{\pi(z^1) + \pi(z^2)} \Big),
\end{aligned} \tag{32}$$

where J is the total number of different probable instances for z. Note that the drawing mechanism we used for z^1 and z^2 in the description of Lemma 2 implies a probability of $\frac{\pi(z^1)+\pi(z^2)}{J}$ for drawing z^1 and z^2 because the probability of drawing z^1 then z^2 is $\pi(z^1) \times 1/J$ and the probability of drawing z^2 followed by

z^1 is $\pi(z^2) \times 1/J$. Therefore, it follows that

$$
\begin{aligned}
V_t(x_t, z_t) = & \max_{a_t, a(z) \text{ feasible}} w(x_t, z_t, a_t) \\
& + \mathop{\mathbb{E}}_{z^1, z^2} \Big[\frac{\pi(z^1)\big(w(x_{t+1}, a(z^1), z^1) + V_{t+2}(\Gamma(x_{t+1}, a(z^1), z^1))\big)}{\pi(z^1) + \pi(z^2)} \\
& + \frac{\pi(z^2)\big(w(x_{t+1}, a(z^2), z^2) + V_{t+2}(\Gamma(x_{t+1}, a(z^2), z^2))\big)}{\pi(z^1) + \pi(z^2)} \Big] \\
= & \max_{a_t, a(z) \text{ feasible}} w(x_t, z_t, a_t) \\
& + \mathop{\mathbb{E}}_{z^1, z^2} \Big[\pi^1 . \big(w(x_{t+1}, a(z^1), z^1) + V_{t+2}(\Gamma(x_{t+1}, a(z^1), z^1))\big) \\
& + \pi^2 . \big(w(x_{t+1}, a(z^2), z^2) + V_{t+2}(\Gamma(x_{t+1}, a(z^2), z^2))\big) \Big] \\
\leq & \mathop{\mathbb{E}}_{z^1, z^2} \Big[\max_{a_t, a^1, a^2 \text{ feasible}} w(x_t, z_t, a_t) + \pi^1 . \big(w(x_{t+1}, a^1, z^1) + V_{t+2}(\Gamma(x_{t+1}, a^1, z^1))\big) \\
& + \pi^2 . \big(w(x_{t+1}, a^2, z^2) + V_{t+2}(\Gamma(x_{t+1}, a^2, z^2))\big) \Big]
\end{aligned}
\tag{33}
$$

□

Lemma 3. *Define a stochastic dynamic program D1 and draws z^1 and z^2 as in Lemma 2. The following inequality holds:*

$$
\begin{aligned}
V_t(x_t, z_t) = & \max_{a_t \in A(x_t, z_t)} w(x_t, z_t, a_t) + \mathop{\mathbb{E}}_{z_{t+1}} [V_{t+1}(\Gamma(x_t, a_t, z_t), z_{t+1})] \\
\leq & \mathop{\mathbb{E}}_{z^1, z^2} \big[w(x_t, z_t, a_t(z^1, z^2)) + w(x_{t+1}, a^1(z^1, z^2), z^1) + V_{t+2}(\Gamma(x_{t+1}, a^1(z^1, z^2), z^1)) \big]
\end{aligned}
\tag{34}
$$

where $a^t(z^1, z^2)$ and $a_1(z^1, z^2)$ are given by:

$$
\begin{aligned}
& [a_t(z^1, z^2), a^1(z^1, z^2), a^2(z^1, z^2)] \\
\triangleq & \arg \max_{a_t, a^1, a^2} w(x_t, z_t, a_t) + \pi^1 . \big(w(x_{t+1}, a^1, z^1) + V_{t+2}(\Gamma(x_{t+1}, a^1, z^1))\big) \\
& + \pi^2 . \big(w(x_{t+1}, a^2, z^2) + V_{t+2}(\Gamma(x_{t+1}, a^2, z^2))\big) \big] \\
& \textit{s.t.}\ \ a_t \in A(x_t, z_t),\ \ x_{t+1} = \Gamma(x_t, a_t, z_t),\ \ a^1, a^2 \in A(x_t, z_t).
\end{aligned}
\tag{35}
$$

Proof.

$$
\begin{aligned}
&\mathop{\mathbb{E}}_{z^1,z^2}\left[w(x_t,z_t,a_t(z^1,z^2))+w(x_{t+1},a^1(z^1,z^2),z^1)+V_{t+2}(\Gamma(x_{t+1},a^1(z^1,z^2),z^1))\right]\\
=&\mathop{\mathbb{E}}_{z^2}\Big[\sum_{z^1}\pi(z^1)\big(w(x_t,z_t,a_t(z^1,z^2))+w(x_{t+1},a^1(z^1,z^2),z^1)+V_{t+2}(\Gamma(x_{t+1},a^1(z^1,z^2),z^1))\big)\Big]\\
=&\frac{1}{J}\sum_{z^2}\sum_{z^1}\pi(z^1)\big(w(x_t,z_t,a_t(z^1,z^2))+w(x_{t+1},a^1(z^1,z^2),z^1)+V_{t+2}(\Gamma(x_{t+1},a^1(z^1,z^2),z^1))\big) \qquad (36)\\
=&\sum_{z^1,z^2}\Big(\frac{\pi(z^1)}{J}\big(w(x_t,z_t,a_t(z^1,z^2))+w(x_{t+1},a^1(z^1,z^2),z^1)+V_{t+2}(\Gamma(x_{t+1},a^1(z^1,z^2),z^1))\big)\\
&\quad+\frac{\pi(z^2)}{J}\big(w(x_t,z_t,a_t(z^1,z^2))+w(x_{t+1},a^2(z^1,z^2),z^1)+V_{t+2}(\Gamma(x_{t+1},a^2(z^1,z^2),z^1))\big)\Big)
\end{aligned}
$$

where the rightmost equality follows by rearranging the sums and from the fact that the sums are taken over all possible realizations of z^1 and z^2, where z^1 and z^2 share the same set of possible values. Finally, by the definitions of $a_t(z^1,z^2)$, $a^1(z^1,z^2)$, and $a^2(z^1,z^2)$ according to (35), the right-hand side of (36) is equal to the right-hand side of (32). Therefore, the rest of the proof follows from the proof of Lemma 2. □

We now prove Proposition 3 using Lemmas 1-3. First, we use Lemma 1 and 2 to show that $\mathbb{E}[\tilde{V}]$ is an upper bound on the optimal revenue V_0^*. A similar reasoning shows that $\mathbb{E}[\sum_d e_{d1}u_dw_d]$ (i.e. the realized revenue of the first sequence) also provides an upper bound on V_0^*; using Lemma 3 instead of Lemma 2.

Define a variable $\gamma_{td}\in\{0,1\}$ determining whether request d belongs to period t. Note that each request belongs to one and only one period. Rewrite problem (9) as:

$$
\begin{aligned}
\tilde{V}&=\max_{\{u_d\}\text{feasible}}\ \sum_s\pi_s\sum_d e_{ds}u_dw_d\\
&=\max_{\{u_d\}\text{feasible}}\ \sum_t\sum_s\pi_s\sum_d\gamma_{td}e_{ds}u_dw_d \qquad (37)\\
&=\max_{\{u_d\}\text{feasible}}\ \sum_{t=1}^{t_M-1}\sum_s\pi_s\sum_d\gamma_{td}e_{ds}u_dw_d+\sum_{t=t_M}^{T}\sum_s\pi_s\sum_d\gamma_{td}e_{ds}u_dw_d.
\end{aligned}
$$

Define the stochastic dynamic program DP1 as follows: for periods 1 to t_M-1, the decision sequences are the same as those drawn in the Simulation Procedure. For periods t_M to T the draws are ignored, requests arrive according to their regular distributions π_{jt}. Let $t'=t_M-1$ and $t''=T+1$, using Lemma 1 we infer:

$$
\begin{aligned}
&\max_{u_d,u_t(\{j(t)\})\text{ feasible}}\ \sum_{t=1}^{t_M-1}\sum_s\pi_s\sum_d\gamma_{td}e_{ds}u_dw_d+\sum_{\{j(t)\}}P(\{j(t)\})\sum_{t=t_M}^{T}u_t(\{j(t)\})w_t(j(t))\\
\geq&\max_{u_d\text{ feasible}}\sum_s\pi_s\Big(\sum_{t=1}^{t_M-1}\sum_d\gamma_{td}e_{ds}u_dw_d+V_{t_M}(\mathbf{C}-\sum_{t=1}^{t_M-1}\sum_d\gamma_{td}e_{ds}u_d\mathbf{L}_d)\Big)
\end{aligned}
\qquad (38)
$$

where $\{j(t)\}$ is a random realization of the service requests received from period t_M onward and $P(\{j(t)\})$ is the probability of this realization. The decisions $u_t(\{j(t)\})$ are made based on the full realization when it is known at time t_M. To use Lemma 2, let the stochastic dynamic program D1 follow the decision sequences up to period $t_M - 1$, then for period t_M let the random sequence $\{j(t)\}$ correspond to z_{t+1} in Lemma 2 and let the respective decisions $u_t(\{j(t)\})$ correspond to the action $a(z_{t+1})$ in (33). Lemma 2 implies that reducing the set of possible future sequences $\{j(t)\}$ to two sequences (using the described drawing mechanism) increases the expected value function:

$$\begin{aligned}
\mathbb{E}[\tilde{V}] &= \mathbb{E}\Big[\max_{\{u_d\}\text{feasible}} \sum_{t=1}^{t_M-1}\sum_s \pi_s \sum_d \gamma_{td} e_{ds} u_d w_d + \sum_{t=t_M}^{T}\sum_s \pi_s \sum_d \gamma_{td} e_{ds} u_d w_d\Big] \\
&\geq \mathbb{E}\Big[\max_{u_d, u_t(\{j(t)\})\text{ feasible}} \sum_{t=1}^{t_M-1}\sum_s \pi_s \sum_d \gamma_{td} e_{ds} u_d w_d + \sum_{\{j(t)\}} P(\{j(t)\}) \sum_{t=t_M}^{T} u_t(\{j(t)\}) w_t(j(t))\Big] \qquad (39) \\
&\geq \mathbb{E}\Big[\max_{u_d\text{ feasible}} \sum_s \pi_s \Big(\sum_{t=1}^{t_M-1}\sum_d \gamma_{td} e_{ds} u_d w_d + V_{t_M}(\mathbf{C} - \sum_{t=1}^{t_M-1}\sum_d \gamma_{td} e_{ds} u_d \mathbf{L}_d)\Big)\Big].
\end{aligned}$$

By splitting the right-hand side of the inequality (39) into periods 1 to $T_{M-1} - 1$ and T_{M-1} to $T_M - 1$, we can use a similar reasoning as above to infer that:

$$\mathbb{E}[\tilde{V}] \geq \mathbb{E}\Big[\max_{u_d\text{ feasible}} \sum_s \pi_s \Big(\sum_{t=1}^{t_{M-1}-1}\sum_d \gamma_{td} e_{ds} u_d w_d + V_{t_{M-1}}(\mathbf{C} - \sum_{t=1}^{t_{M-1}-1}\sum_d \gamma_{td} e_{ds} u_d \mathbf{L}_d)\Big)\Big]. \qquad (40)$$

Hence by backward induction and repeated use of Lemma 1 and Lemma 2, we reach the conclusion. □

Proposition 5.

Proof. We prove that $V^S(\{j(t)\})$ is an upper bound on $\check{V}_1(\mathbf{C}|\{j(t)\})$, and the statement follows from the fact that $\check{V}_1(\mathbf{C}) = \mathbb{E}_{\{j(t)\}}\check{V}_1(\mathbf{C}|\{j(t)\})$ is an upper bound on $V_1(\mathbf{C})$. In this proof we focus on the adjusted problem below. The proof extends to the original problem (18) because the optimal solution to the adjusted problem is a feasible solution for (18).

$$[SEMI] \quad V^S(\{j(t)\}) \triangleq \max_{\{\eta_t(n),\gamma_t(n)\}} \Big[\sum_{t=1}^{T}\sum_{n=1}^{N} \gamma_t(n)\Big]$$

$$\text{s.t.} \quad \sum_{n=1}^{N} \gamma_t(n) \leq W_m^t \ \ \forall\, t \ , \ \sum_{t=1}^{T}\sum_{n=1}^{N} \eta_t(n) L_{j(t)t}(n) \leq C \qquad (41a)$$

$$\sum_{n=1}^{N} \eta_t(n) \leq 1 \ \ \forall\, t \ , \ \ 0 \leq \eta_t(n) \ \ \forall\, n, t \qquad (41b)$$

$$\gamma_t(n) \leq \eta_t(n) f_{j(t)}(n)\big(\eta_t(n)\big) \ \ \forall\, n, t \qquad (41c)$$

Assume that for a given realization of service requests $\{j(t)\}$ we are able to solve (17) and act optimally (resulting in an average total revenue of $\check{V}_1(\mathbf{C}|\{j(t)\})$). In this optimal policy, let $\hat{\eta}_t(n)$ denote the probability that option n is chosen in period t, and let $\hat{\gamma}_t(n)$ denote the expected revenue made from option n in period t. By definition, $\check{V}_1(\mathbf{C}|\{j(t)\}) = \sum_{t=1}^{T}\sum_{n=1}^{N}\hat{\gamma}_t(n)$. If we can show that $\hat{\eta}_t(n)$ and $\hat{\gamma}_t(n)$ are feasible for problem $[Semi]$, it follows that $V^S(\{j(t)\})$ is at least as large as $\check{V}_1(\mathbf{C}|\{j(t)\})$, yielding the conclusion we seek. Since the optimal policy is a feasible policy, the constraints in (41a) and (41b) are satisfied by definition. Constraint (41c) requires further discussion.

To show that $\hat{\eta}_t(n)$ and $\hat{\gamma}_t(n)$ satisfy (41c), we show that the realized average revenue from option n in period t, under the optimal policy, cannot exceed $\eta_t(n)f_{j(t)}(n)\big(\eta_t(n)\big)$; where $\eta_t(n)$ is the probability of option n being chosen under the implemented policy. First note that in the pairs $(y_j^k(n), w_j^k(n))$ that were used to define $f_j(n)\big(\eta\big)$, the probability of option n being chosen given price $w_t^k(n)$ is overestimated because $y_j^k(n)$ is determined assuming no other options are offered. The other options would reduce the probability of option n being chosen because some customer types might switch to another option. Effectively, the realized choice probability will be $\hat{y}_j^k(n) \le y_j^k(n)$ depending on the frequency of prices for other options. Assume that we know the realized $\hat{y}_j^k(n)$ in the optimal policy. Define the function $g_j(n)\big(\eta\big)$ as the smallest non-increasing concave function that satisfies $g_j(n)(\hat{y}_j^k(n)) \ge p_j^k(n) \ \forall k \in \{1,...,K\}$. It is easy to see that $g_j(n)\big(\eta\big) \le f_j(n)\big(\eta\big) \ \forall \eta \in [0,1]$, hence if the constraint (41c) holds by replacing $f_j(n)$ with $g_j(n)$, then it will also hold for $f_j(n)$. In the optimal policy, an average choice probability of $\hat{\eta}_t(n)$ is attained as a mix of the $\hat{y}_t^k(n)$s (because different price points are offered depending on the state at time t). Hence the revenue $\gamma_t(n)$ received from option n will be a weighted average of the revenues at the points $\hat{y}_t^k(n)$. The function $g_j(n)\big(\eta\big)$ is concave non-increasing, thus the function $\eta g_j(n)\big(\eta\big)$ is concave for $\eta \ge 0$. Therefore, the inequality $\hat{\gamma}_t(n) \le \hat{\eta}_t(n)g_j(n)\big(\hat{\eta}_t(n)\big)$ follows from Jenson's inequality. □

B Numerical Setup

In our numerical tests, the predicted user traffic pattern for a type j application is drawn from a query from Google Trends for a 100-day period ($R = 100$). These queries are used to form $\mathbf{L}_{jt}$, appropriately scaled according to a capacity of $c = 100$. Given $\mathbf{L}_{jt}$, we describe the functions $\mathbf{O}_n(\mathbf{L})$ and $w_{kn}(\mathbf{L})$ used to generate the menu options and willingness-to-pay for the test problems in §3.3. For the test problems in §2.3, we use the given patterns $\mathbf{L}_{jt}$ and let $w_j \doteq w_{k'1}(\mathbf{L}_{jt})$, where k' is a randomly chosen customer type in $[1,...,K]$.

Menu Options $\mathbf{O}_n$.

(i) The first option is just to support the predicted user traffic $\mathbf{L}$. $\mathbf{O}_1(\mathbf{L}) = \mathbf{L}$.

(ii) The second offer is to support 1.2 times the predicted user traffic. $\mathbf{O}_2(\mathbf{L}) = \mathbf{1.2L}$.

(iii) The third offer is to support the maximum user traffic over the entire horizon. $\mathbf{O}_3(\mathbf{L}) = [\ell^m, \ell^m, ..., \ell^m]$, where $\ell^m = \max_r\{\ell_r\}$.

Willingness to Pay w_{kn}.

We consider three categories of customer characteristics and preferences. Each customer type k takes one type from each category, creating a total of 12 different customer types. The function $w_{kn}(\mathbf{L})$ is a product of three elements, each pertaining to one of the mentioned classes: $w_{kn}(\mathbf{L}) = S\omega_k(\mathbf{L})\rho_{kn}\upsilon_k$. The parameter S is fixed at \$0.5 to reflect the current cloud market: it approximately costs \$50 to serve an average of 100 simultaneous users during a day according to a price estimate by Microsoft Azure and an estimate of system requirements by Kcura.

(I) The first category of customer characteristics determines whether the customer values scarce resources higher, or values all resources equally.

Type 1: Values all resources equally:

$$\omega_k(\mathbf{L}) = \frac{\sum_r \ell_r}{R}.$$

Type 2: Values scarce resources higher:

$$\omega_k(\mathbf{L}) = \frac{\sum_r \lambda_r \ell_r}{\sum_r \lambda_r}$$

(II) The second category determines whether the customer has a high, medium, or low sensitivity to the risk of insufficient capacity. The willingness to pay for the baseline option (option 1) is set equal for the three types, but the more sensitive types are willing to pay more for the improved service guarantees relative to less sensitive types.

Type 1: Medium variance: $\rho_{k1} = 1$, $\rho_{k2} = 1.2$, $\rho_{k3} = 2$.

Type 2: Low variance: $\rho_{k1} = 1$, $\rho_{k2} = 1.1$, $\rho_{k3} = 1.5$.

Type 3: High variance: $\rho_{k1} = 1$, $\rho_{k2} = 1.5$, $\rho_{k3} = 3$.

(III) The third category determines whether the customer is of "high valuation" or "low valuation" type.

Type 1: high valuation: $\upsilon_k = 1$.

Type 2: low valuation: $\upsilon_k = 0.5$.

Customer Arrival Probabilities $\pi_{\mathbf{jt}}$.

We set the arrival probabilities randomly, such that $\sum_j \pi_{jt} = 1$ for all t, then scale down all probabilities equally so that $\sum_{rjt} \phi_{jt}\ell_{rjt}/cR = \kappa$ (i.e. the total resources requested over time is κ times the capacity of the provider). The reported results are averages of cases with $\kappa \in [1.2, 1.5, 2, 3]$, we did not observe significant difference between results of different values of κ.