Proofs for "Logarithmic Regret in Multisecretary and Online Linear Programs with Continuous Valuations"

Robert L. Bray

Kellogg School of Management, Northwestern University

May 11, 2024

Lemma 1 Proof. Assumption 6 implies that the event $\Delta_1(y) = 0$ and $a_1 \neq 0$ has measure zero, for y sufficiently close to y^β_∞. Hence, $\Delta_1(y)^+$ is almost surely differentiable in y, which means that

$$\begin{aligned}\tfrac{\partial}{\partial y}\Lambda^b_\infty(y) &= \tfrac{\partial}{\partial y}\left(b'y + \mathrm{E}(\Delta_1(y)^+)\right)\\ &= b + \mathrm{E}\left(\tfrac{\partial}{\partial y}\Delta_1(y)^+\right)\\ &= b - \mathrm{E}(\mathbb{1}\{\Delta_1(y) > 0\}a_1).\end{aligned}$$

Note we can commute the expectation and differentiation because a_1 is bounded. Combining the derivative above with Assumption 6 and the convexity of Λ^b_∞ implies the result. □

Lemma 2 Proof. Assumption 5 and Lemma 1 imply that (i) $\dot\Lambda^\beta_\infty(y^\beta_\infty) = 0$, (ii) $\ddot\Lambda_\infty(y^\beta_\infty)$ is nonsingular, and (iii) $\dot\Lambda^b_\infty(y)$ is continuously differentiable in y near y^β_∞. Further, $\dot\Lambda^b_\infty(y)$ is continuously differentiable in b, since $\frac{\partial}{\partial b}\dot\Lambda^b_\infty(y) = I$ (see the proof of Lemma 1). Accordingly, the implicit function theorem establishes that each b in a neighborhood of β has a corresponding shadow price vector y^b_∞ that has continuous derivative $\frac{\partial}{\partial b}y^b_\infty = -\frac{\partial}{\partial y}\dot\Lambda^b_\infty(y)^{-1}\frac{\partial}{\partial b}\dot\Lambda^b_\infty(y)|_{y=y^b_\infty} = -\ddot\Lambda^b_\infty(y^b_\infty)^{-1}$. Further, y^b_∞ must be the unique minimizer of Λ^b_∞ for b near β, because $\dot\Lambda^b_\infty(y^b_\infty) = 0$ and $\ddot\Lambda_\infty(y)$ is positive definite for y near y^β_∞. □

Corollary 1 Proof. This follows immediately from Theorem 3. □

Proposition 3 Proof. I will first establish that Σ^b is continuous and full rank for all b in a neighborhood of β. Lemmas 1 and 2 imply the continuity, and Lemma 1 implies that Σ^b is full rank if $\mathrm{Cov}(\mathbb{1}\{\Delta_1(y^b_\infty) > 0\}a_1)$ is full rank. If this latter matrix were not full rank, then there would be some $\gamma \neq 0$ that almost surely satisfies $\mathbb{1}\{\Delta_1(y^b_\infty) > 0\}a_1'\gamma = \mathrm{E}(\mathbb{1}\{\Delta_1(y^b_\infty) > 0\}a_1'\gamma)$, which would imply that either (i) $\Delta_1(y^b_\infty) > 0$, almost surely, or (ii) $\mathbb{1}\{\Delta_1(y^b_\infty) > 0\}a_1'\gamma = 0$, almost surely. The former case violates Assumption 6 because it implies that $\mathrm{E}(\mathbb{1}\{\Delta_1(y+dy) > 0\}a_1'\gamma) = \mathrm{E}(\mathbb{1}\{\Delta_1(y) > 0\}a_1'\gamma)$ for $dy \leq 0$, and the latter case violates Assumption 6 because it implies that $\mathrm{E}(\mathbb{1}\{\Delta_1(y+dy) > 0\}a_1'\gamma) = \mathrm{E}(\mathbb{1}\{\Delta_1(y) > 0\}a_1'\gamma)$ for $dy > 0$.

The fact that $\sqrt{t}(y^b_t - y^b_\infty) \xrightarrow{d} \mathcal{N}(0, \Sigma^b)$ follows directly from theorem 2.13 of Kosorok (2008), so it will suffice to show that the conditions of this theorem hold. To use follow Kosorok's notation, define functions

$$
\begin{aligned}
m_y(u_1, a_1) &\equiv b'y + \Delta_1(y)^+, \\
\dot{m}(a_1) &\equiv \|b\| + \|a_1\|, \\
\text{and} \quad \dot{m}_\infty(u_1, a_1) &\equiv b - \mathbb{1}\{\Delta_1(y^b_\infty) > 0\}a_1.
\end{aligned}
$$

First, the Hessian matrix of $\mathrm{E}(m_y(u_1, a_1))$ at $y = y^b_\infty$ is $\ddot{\Lambda}_\infty(y^b_\infty)$, which is non-singular when b is sufficiently close to β, by Lemmas 1 and 2. Second, Assumption 4 establishes that $\mathrm{E}(\dot{m}(a_1)^2)$ and $\mathrm{E}(\|\dot{m}_\infty(u_1, a_1)\|^2)$ are finite. Third, functions m_y and $\dot{m}$ satisfy condition (2.18) of Kosorok (2008):

$$
\begin{aligned}
|m_y(u_1, a_1) - m_z(u_1, a_1)| &= b'y + \Delta_1(y)^+ - b'z - \Delta_1(z)^+ \\
&\leq (\|b\| + \|a_1\|)\|y - z\| \\
&= \dot{m}(a_1)\|y - z\|.
\end{aligned}
$$

Fourth, Assumption 6 ensures that functions m_y and $\dot{m}_\infty$ satisfy condition (2.19) of Kosorok (2008):

$$\begin{aligned}
&\mathrm{E}\left(\left(m_y(u_1,a_1) - m_{y^b_\infty}(u_1,a_1) - \dot{m}_\infty(u_1,a_1)'(y - y^b_\infty)\right)^2\right) \\
&= \mathrm{E}\left(\left(\Delta_1(y)^+ - \Delta_1(y^b_\infty)^+ + \mathbb{1}\{\Delta_1(y^b_\infty) > 0\}a_1'(y - y^b_\infty)\right)^2\right) \\
&= \mathrm{E}\left(\Delta_1(y)^2 |\mathbb{1}\{\Delta_1(y) > 0\} - \mathbb{1}\{\Delta_1(y^b_\infty) > 0\}|\right) \\
&\leq \mathrm{E}\left((a_1'y - a_1'y^b_\infty)^2 |\mathbb{1}\{\Delta_1(y) > 0\} - \mathbb{1}\{\Delta_1(y^b_\infty) > 0\}|\right) \\
&= \|y - y^b_\infty\|^2 \,\mathrm{E}\left(\|\mathbb{1}\{\Delta_1(y) > 0\}a_1 - \mathbb{1}\{\Delta_1(y^b_\infty) > 0\}a_1\|^2\right) \\
&\leq \|\alpha\| \|y - y^b_\infty\|^2 \,\mathrm{E}\left(\mathbb{1}\{\Delta_1(y \wedge y^b_\infty) > 0\}a_1 - \mathbb{1}\{\Delta_1(y \vee y^b_\infty) > 0\}a_1\right) \\
&\leq \|\alpha\| \|y - y^b_\infty\|^2 O(\|y - y^b_\infty\|) \\
&= o(\|y - y^b_\infty\|).
\end{aligned}$$

Finally, Proposition 4 establishes that $\|y^b_t - y^b_\infty\| \xrightarrow{p} 0$. □

Proposition 4 Proof. Since Proposition 6 establishes that $\mathrm{E}(\sup_{b \in B_\delta(\beta)} \mathbb{1}\{y^b_t \notin B_\epsilon(y^b_\infty)\}\|y^b_t - y^b_\infty\|^2) = o(1/t)$, it will suffice to show that $\mathrm{E}(\sup_{b \in B_\delta(\beta)} \mathbb{1}\{y^b_t \in B_\epsilon(y^b_\infty)\}\|y^b_t - y^b_\infty\|^2) = O(1/t)$, for sufficiently small $\epsilon > 0$. I will establish this result with Theorems 2.14.2 and 2.14.5 of van der Vaart and Wellner (1996). However, translating the problem into van der Vaart and Wellner's empirical processes framework will take some effort. First, I bound the magnitude of $y^b_t - y^b_\infty$ in terms of the magnitude of $\dot{\Lambda}^b_\infty(\hat{y}^b_t) - \dot{\Lambda}^b_t(\hat{y}^b_t)$, where $\hat{y}^b_t \equiv (y^b_t + y^b_\infty)/2$. Since $\hat{y}^b_t$ lies between the minimizers of Λ^b_∞ and Λ^b_t, the vector $\hat{y}^b_t - y^b_\infty$ projects positively onto gradient $\dot{\Lambda}^b_\infty(\hat{y}^b_t)$ and projects negatively onto subgradient $\dot{\Lambda}^b_t(\hat{y}^b_t)$. I use this fact to show that $(\hat{y}^b_t - y^b_\infty)'(\dot{\Lambda}^b_\infty(\hat{y}^b_t) - \dot{\Lambda}^b_t(\hat{y}^b_t))$ is larger than some fixed multiple of $\|\hat{y}^b_t - y^b_\infty\|^2$, which indicates that $\|\dot{\Lambda}^b_\infty(\hat{y}^b_t) - \dot{\Lambda}^b_t(\hat{y}^b_t)\|$ is larger than some fixed multiple of $\|y^b_t - y^b_\infty\|$. This, in turn, implies that the expectation of the maximum of $\|\dot{\Lambda}^b_t(y) - \dot{\Lambda}^b_\infty(y)\|^2$, across y in some small ball of y^β_∞, is larger than some fixed multiple of the expectation of $\mathbb{1}\{\|y^b_t - y^b_\infty\| \leq \epsilon\}\|y^b_t - y^b_\infty\|^2$. And bounding the expectation of the maximum of $\|\dot{\Lambda}^b_t(y) - \dot{\Lambda}^b_\infty(y)\|^2$ is a classic empirical processes problem.

Now, let's get to the proof. First, Lemma 2 establishes that we can choose δ small enough so that $y^b_\infty \in B_\epsilon(y^\beta_\infty)$ for all $b \in B_\delta(\beta)$, in which case $y^b_t \in B_\epsilon(y^b_\infty)$ implies $y^b_t \in B_{2\epsilon}(y^\beta_\infty)$, which in turn implies $\hat{y}^b_t \in B_{3\epsilon/2}(y^b_\infty)$, where $\hat{y}^b_t \equiv (y^b_t + y^b_\infty)/2$.

Second, let σ^b_m denote the smallest singular value of $\ddot{\Lambda}_\infty(y^b_\infty)$. Lemmas 1 and 2 imply that we can set δ small enough so that for all $b \in B_\delta(\beta)$ we have $\sigma_m^b \geq \sigma_m^\beta/2$, and hence

$$(\hat{y}_t^b - y_\infty^b)'\ddot{\Lambda}_\infty(y_\infty^b)(\hat{y}_t^b - y_\infty^b) \geq \sigma_m^\beta \|\hat{y}_t^b - y_\infty^b\|^2/2.$$

Next, note that $\dot{\Lambda}_\infty^b(y_\infty^b) = 0$ implies

$$\begin{aligned}\dot{\Lambda}_\infty^b(\hat{y}_t^b) &= \dot{\Lambda}_\infty^b(\hat{y}_t^b) - \dot{\Lambda}_\infty^b(y_\infty^b) \\ &= \ddot{\Lambda}_\infty(y_\infty^b)(\hat{y}_t^b - y_\infty^b) + o(\|\hat{y}_t^b - y_\infty^b\|),\end{aligned}$$

where the little-o term holds uniformly across $b \in B_\delta(\beta)$. Accordingly, we can set ϵ small enough so that $y_t^b \in B_{2\epsilon}(y_\infty^\beta)$ implies

$$\|\dot{\Lambda}_\infty^b(\hat{y}_t^b) - \ddot{\Lambda}_\infty(y_\infty^b)(\hat{y}_t^b - y_\infty^b)\| \leq \sigma_m^\beta \|\hat{y}_t^b - y_\infty^b\|/4,$$

for all $b \in B_\delta(\beta)$. Now combining these last two results yields the following, for $y_t^b \in B_{2\epsilon}(y_\infty^\beta)$:

$$\begin{aligned}(\hat{y}_t^b - &y_\infty^b)'\dot{\Lambda}_\infty^b(\hat{y}_t^b) \\ &= (\hat{y}_t^b - y_\infty^b)'\ddot{\Lambda}_\infty(y_\infty^b)(\hat{y}_t^b - y_\infty^b) \\ &\quad + (\hat{y}_t^b - y_\infty^b)'\big(\dot{\Lambda}_\infty^b(\hat{y}_t^b) - \ddot{\Lambda}_\infty(y_\infty^b)(\hat{y}_t^b - y_\infty^b)\big) \\ &\geq \sigma_m^\beta \|\hat{y}_t^b - y_\infty^b\|^2/2 - \|\hat{y}_t^b - y_\infty^b\|\|\dot{\Lambda}_\infty^b(\hat{y}_t^b) - \ddot{\Lambda}_\infty(y_\infty^b)(\hat{y}_t^b - y_\infty^b)\| \\ &\geq \sigma_m^\beta \|\hat{y}_t^b - y_\infty^b\|^2/4.\end{aligned}$$

And combining this with $(\hat{y}_t^b - y_\infty^b)'\dot{\Lambda}_t^b(\hat{y}_t^b) = (y_t^b - \hat{y}_t^b)'\dot{\Lambda}_t^b(\hat{y}_t^b) \leq 0$, which we get from Lemma 1, yields the following, for $y_t^b \in B_{2\epsilon}(y_\infty^\beta)$:

$$\begin{aligned}\|\hat{y}_t^b - y_\infty^b\|&\|\dot{\Lambda}_\infty^b(\hat{y}_t^b) - \dot{\Lambda}_t^b(\hat{y}_t^b)\| \\ &\geq (\hat{y}_t^b - y_\infty^b)'(\dot{\Lambda}_\infty^b(\hat{y}_t^b) - \dot{\Lambda}_t^b(\hat{y}_t^b)) \\ &\geq \sigma_m^\beta \|\hat{y}_t^b - y_\infty^b\|^2/4.\end{aligned}$$

Hence, $y_t^b \in B_{2\epsilon}(y_\infty^\beta)$ implies

$$\|\dot{\Lambda}_\infty^b(\hat{y}_t^b) - \dot{\Lambda}_t^b(\hat{y}_t^b)\| \geq \sigma_m^\beta \|\hat{y}_t^b - y_\infty^b\|/4 = \sigma_m^\beta \|y_t^b - y_\infty^b\|/8.$$

And thus, we have

$$
\begin{aligned}
&\mathrm{E}\Big(\sup_{b \in B_\delta(\beta)} \mathbb{1}\{y_t^b \in B_\epsilon(y_\infty^b)\} \|y_t^b - y_\infty^b\|^2 \Big) \\
&\quad \le \mathrm{E}\Big(\sup_{b \in B_\delta(\beta)} \mathbb{1}\{y_t^b \in B_{2\epsilon}(y_\infty^\beta)\} \|y_t^b - y_\infty^b\|^2 \Big) \\
&\quad \le (8/\sigma_m^\beta)^2 \, \mathrm{E}\Big(\sup_{b \in B_\delta(\beta)} \mathbb{1}\{y_t^b \in B_{2\epsilon}(y_\infty^\beta)\} \|\dot\Lambda_t^b(\hat y_t^b) - \dot\Lambda_\infty^b(\hat y_t^b)\|^2 \Big) \\
&\quad \le (8/\sigma_m^\beta)^2 \, \mathrm{E}\Big(\sup_{b \in B_\delta(\beta)} \sup_{y \in B_{2\epsilon}(y_\infty^\beta)} \|\dot\Lambda_t^b(y) - \dot\Lambda_\infty^b(y)\|^2 \Big) \\
&\quad = (8/\sigma_m^\beta)^2 \, \mathrm{E}\Big(\sup_{y \in B_{2\epsilon}(y_\infty^\beta)} \|\dot\Lambda_t^\beta(y) - \dot\Lambda_\infty^\beta(y)\|^2 \Big),
\end{aligned}
$$

where the last line holds because $\dot\Lambda_t^b - \dot\Lambda_\infty^b$ is independent of b. Finally, Lemma 5 establishes that the expectation in the last line is less than C/t for some universal constant $C > 0$. □

Proposition 5 Proof. This follows immediately from Lemma 2. □

Corollary 2 Proof. This follows from Proposition 4 and Lemma 2. □

Proposition 6 Proof. The proof hinges on two key results. The first result is that there exists $\delta, C > 0$ such that

$$
\Pr\big(\sup_{b \in B_\delta(\beta)} \|y_t^b - y_\infty^b\| > \epsilon \big) \le 4m^2 \exp(-C\epsilon^2 t), \tag{1}
$$

for all $t \in \mathbb{N}$ and sufficiently small $\epsilon > 0$. The second result is that for all sufficiently large $\gamma > 0$ there exists $\delta, C > 0$ such that

$$
\mathrm{E}\big(\sup_{b \in B_\delta(\beta)} \mathbb{1}\{y_t^b \notin B_{\gamma^{1/p}}(0)\} \|y_t^b\|^p \big) \le \exp(-Ct), \tag{2}
$$

for all sufficiently large t.

The $p = 0$ case follows immediately from the line (1). Deriving the $p > 0$ case from lines (1) and (2) will take a bit more work. To that end, choose γ large enough so that $\gamma \ge \sup_{b \in B_\delta(\beta)} \|y_\infty^b\|^p$, and hence $\|y_t^b - y_\infty^b\| \le \|y_t^b\| + \gamma^{1/p}$ (Lemma 2 establishes that this is possible). And with this, lines (1) and (2) imply that we can choose $C > 0$ so that we have the following for all sufficiently small

ϵ and large t:

$$
\begin{aligned}
&\mathrm{E}\big(\sup_{b\in B_\delta(\beta)} \mathbb{1}\{y_t^b \notin B_\epsilon(y_\infty^b)\}\|y_t^b - y_\infty^b\|^p\big)\\
&\quad\le \mathrm{E}\big(\sup_{b\in B_\delta(\beta)} \mathbb{1}\{y_t^b \notin B_\epsilon(y_\infty^b)\}\mathbb{1}\{y_t^b \notin B_{\gamma^{1/p}}(0)\}(\|y_t^b\| + \gamma^{1/p})^p\big)\\
&\qquad + \mathrm{E}\big(\sup_{b\in B_\delta(\beta)} \mathbb{1}\{y_t^b \notin B_\epsilon(y_\infty^b)\}\mathbb{1}\{y_t^b \in B_{\gamma^{1/p}}(0)\}(\|y_t^b\| + \gamma^{1/p})^p\big)\\
&\quad\le \mathrm{E}\big(\sup_{b\in B_\delta(\beta)} \mathbb{1}\{y_t^b \notin B_\epsilon(y_\infty^b)\}\mathbb{1}\{y_t^b \notin B_{\gamma^{1/p}}(0)\}2^p\|y_t^b\|^p\big)\\
&\qquad + \mathrm{E}\big(\sup_{b\in B_\delta(\beta)} \mathbb{1}\{y_t^b \notin B_\epsilon(y_\infty^b)\}\mathbb{1}\{y_t^b \in B_{\gamma^{1/p}}(0)\}2^p\gamma\big)\\
&\quad\le 2^p\,\mathrm{E}\big(\sup_{b\in B_\delta(\beta)} \mathbb{1}\{y_t^b \notin B_{\gamma^{1/p}}(0)\}\|y_t^b\|^p\big) + 2^p\gamma\,\Pr\big(\sup_{b\in B_\delta(\beta)} \|y_t^b - y_\infty^b\| > \epsilon\big)\\
&\quad\le 2^p\exp(-Ct) + 2^{p+2}\gamma m^2\exp(-C\epsilon^2 t).
\end{aligned}
$$

The inequality above establishes the $p > 0$ case. Hence, proving lines (1) and (2) will complete the argument.

Before getting into the math, let me roughly sketch the proof of line (1). The key tool will be Lemma 3, which is our only means for positioning y_t^b. The lemma corresponds to a set of inequalities that describe a small box, which aligns roughly with the orthonormal basis $\{\omega_j^b\}_{i=1}^m$; if these inequalities all hold, then the box is intact, and y_t^b resides inside of it. I will use this result to bound the distance between y_t^b and y_∞^b with the distances between $\dot\Lambda_t^b(y_\infty^b + \eta k\omega_j^b)$ and $\eta k\sigma_j^b\omega_j^b$, for $j \in m$, $k \in \{-1, 1\}$, and $\eta > 0$ (these latter distances being the constraints that ensure the integrity of the box). This reframing simplifies the problem, because $\dot\Lambda_t^b(y_\infty^b + \eta k\omega_j^b)$ is a sum of *i.i.d.* bounded variables. The second part of the proof replaces the $\eta k\sigma_j^b\omega_j^b$ term in our distance measurements with with $\dot\Lambda_\infty^b(y_\infty^b + \eta k\omega_j^b)$. This step is useful because $\dot\Lambda_t^b(y_\infty^b + \eta k\omega_j^b) - \dot\Lambda_\infty^b(y_\infty^b + \eta k\omega_j^b)$ is an empirical process. The final part of the proof invokes a standard empirical process result to establish the desired concentration of measure.

To begin the proof of line (1), note that Lemmas 1 and 2 imply that we can choose $\delta > 0$ and $\epsilon > 0$ small enough to ensure the existence and continuity of $\ddot\Lambda_\infty$ between y_∞^b and $y_\infty^b + \eta k\omega_j^b$, and small enough to ensure that $\sigma_1^b \le 2\sigma_1^\beta$, $\sigma_m^b \ge \sigma_m^\beta/2 > 0$, and $y_\infty^b + \eta k\omega_j^b \ge 0$, for all $j \in [m]$, $k \in \{-1, 1\}$, $b \in B_\delta(\beta)$, and $\eta \equiv \epsilon/(1 + 8\sqrt{m}\sigma_1^\beta/\sigma_m^\beta)$. Now, with these conditions, we can use

Lemma 3 to bound the left-hand side of (1) in terms of more amenable subgradients:

$$
\begin{aligned}
&\Pr\big(\sup_{b\in B_\delta(\beta)} \|y_t^b - y_\infty^b\| > \epsilon\big) \\
&\quad\le \Pr\big(\sup_{b\in B_\delta(\beta)} \|y_t^b - y_\infty^b\| > \eta(1+2\sqrt{m}\sigma_1^b/\sigma_m^b)\big) \\
&\quad\le \Pr\big(\sup_{b\in B_\delta(\beta)} \max_{j\in[m]} \max_{k\in\{-1,1\}} \|\dot{\Lambda}_t^b(y_\infty^b + \eta k\omega_j^b) - \eta k\sigma_j^b\omega_j^b\| - \eta\sigma_m^b/(2\sqrt{m}) > 0\big) \\
&\quad\le \Pr\big(\sup_{b\in B_\delta(\beta)} \max_{j\in[m]} \max_{k\in\{-1,1\}} \|\dot{\Lambda}_t^b(y_\infty^b + \eta k\omega_j^b) - \eta k\sigma_j^b\omega_j^b\| - \eta\sigma_m^\beta/(4\sqrt{m}) > 0\big). \qquad (3)
\end{aligned}
$$

Now I will frame the last expression above as an empirical process by replacing the $\eta k\sigma_j^b\omega_j^b$ term with $\dot{\Lambda}_\infty^b(y_\infty^b + \eta k\omega_j^b)$. To this end, note that the mean value theorem indicates that there exists $\xi \in (0,\eta)$ for which

$$
\begin{aligned}
\dot{\Lambda}_\infty^b(y_\infty^b + \eta k\omega_j^b) &= \dot{\Lambda}_\infty^b(y_\infty^b + \eta k\omega_j^b) - 0 \\
&= \dot{\Lambda}_\infty^b(y_\infty^b + \eta k\omega_j^b) - \dot{\Lambda}_\infty^b(y_\infty^b) \\
&= \eta k\ddot{\Lambda}_\infty(y_\infty^b + \xi k\omega_j^b)\omega_j^b \\
&= \eta k\ddot{\Lambda}_\infty(y_\infty^b)\omega_j^b + \eta k(\ddot{\Lambda}_\infty(y_\infty^b + \xi k\omega_j^b) - \ddot{\Lambda}_\infty(y_\infty^b))\omega_j^b \\
&= \eta k\sigma_j^b\omega_j^b + o(\eta),
\end{aligned}
$$

where the little-o term holds uniformly across $b \in B_\delta(\beta)$. Accordingly, we can set ϵ small enough so that $\sup_{b\in B_\delta(\beta)} \|\dot{\Lambda}_\infty^b(y_\infty^b + \eta k\omega_j^b) - \eta k\sigma_j^b\omega_j^b\| \le \eta\sigma_m^\beta/(8\sqrt{m})$, in which case we have

$$
\begin{aligned}
&\|\dot{\Lambda}_t^b(y_\infty^b + \eta k\omega_j^b) - \eta k\sigma_j^b\omega_j^b\| \\
&\quad\le \|\dot{\Lambda}_t^b(y_\infty^b + \eta k\omega_j^b) - \dot{\Lambda}_\infty^b(y_\infty^b + \eta k\omega_j^b)\| + \|\dot{\Lambda}_\infty^b(y_\infty^b + \eta k\omega_j^b) - \eta k\sigma_j^b\omega_j^b\| \\
&\quad\le \|\dot{\Lambda}_t^b(y_\infty^b + \eta k\omega_j^b) - \dot{\Lambda}_\infty^b(y_\infty^b + \eta k\omega_j^b)\| + \eta\sigma_m^\beta/(8\sqrt{m}).
\end{aligned}
$$

And, finally, combining this with line (3) and the fact that $\dot{\Lambda}_t^b - \dot{\Lambda}_t^b = \dot{\Lambda}_t^\beta - \dot{\Lambda}_t^\beta$ yields the following:

$$
\begin{aligned}
&\Pr\big(\sup_{b\in B_\delta(\beta)} \|y_t^b - y_\infty^b\| > \epsilon\big) \\
&\quad\le \Pr\big(\sup_{b\in B_\delta(\beta)} \max_{j\in[m]} \max_{k\in\{-1,1\}} \|\dot{\Lambda}_t^b(y_\infty^b + \eta k\omega_j^b) - \eta k\sigma_j^b\omega_j^b\| > \eta\sigma_m^\beta/(4\sqrt{m})\big) \\
&\quad\le \Pr\big(\sup_{b\in B_\delta(\beta)} \max_{j\in[m]} \max_{k\in\{-1,1\}} \|\dot{\Lambda}_t^b(y_\infty^b + \eta k\omega_j^b) - \dot{\Lambda}_\infty^b(y_\infty^b + \eta k\omega_j^b)\| > \eta\sigma_m^\beta/(8\sqrt{m})\big) \\
&\quad\le \sum_{i=1}^m\sum_{j=1}^m\sum_{k\in\{-1,1\}} \Pr\big(\sup_{b\in B_\delta(\beta)} |e_i'\dot{\Lambda}_t^b(y_\infty^b + \eta k\omega_j^b) - e_i'\dot{\Lambda}_\infty^b(y_\infty^b + \eta k\omega_j^b)| \ge \eta\sigma_m^\beta/(8m)\big) \\
&\quad\le \sum_{i=1}^m\sum_{j=1}^m\sum_{k\in\{-1,1\}} \Pr\big(\sup_{y\in B_\nu(y_\infty^\beta)} |e_i'\dot{\Lambda}_t^\beta(y) - e_i'\dot{\Lambda}_\infty^\beta(y)| \ge \eta\sigma_m^\beta/(8m)\big),
\end{aligned}
$$

where $\nu > 0$ is a constant that's large enough to ensure that $y_\infty^b + \eta k\omega_j^b \in B_\nu(y_\infty^\beta)$ for all $b \in B_\delta(\beta)$. Finally, Theorem 2.14.9 of van der Vaart and Wellner (1996) implies that this last expression falls exponentially fast in t (see the proof of lemma 5 for confirmation of this theorem's hypothesis). Note that the $\sup_{b\in B_\delta(\beta)}$ prevents us from bounding the probability in the penultimate line with a more standard concentration of measure result.

This establishes line (1), which establishes the $p = 0$ case. I will now prove line (2), assuming $p > 0$. The proof will proceed as follows: First, I will bound the probability that $\|y_t^b\|^p$ exceeds some $\gamma > 0$ with the probability that $e_j'\dot{\Lambda}_t^b(e_j\gamma^{1/p}/\sqrt{m})$ is negative. This latter random variable is easier to work with because it is a sum of *i.i.d.* random variables. Second, I will lower bound $e_j'\dot{\Lambda}_t^b(e_j\gamma^{1/p}/\sqrt{m})$ with a binomial random variable, with success probability $\rho_\gamma \equiv \Pr(u_1 > \eta\gamma^{1/p}/(2\sqrt{m}))$. This characterization will enable me to use the binomial Chernoff bound to establish that $\Pr(\|y_t^b\|^p > \gamma)$ falls exponentially fast in t. And finally, I will integrate over this tail bound to create a corresponding expectation bound.

To begin the proof, I will show that $e_j'y_t^b > \omega$ implies $e_j'\dot{\Lambda}_t^b(\omega e_j) \le 0$, for $\omega \in \mathbb{R}$. To see this, take $e_j' y_t^b \geq \omega$ and $\hat{y} \equiv y_t^b - e_j(e_j' y_t^b - \omega)/2$, and apply Lemma 1:

$$
\begin{aligned}
0 \geq & (y_t^b - \hat{y})' \dot{\Lambda}_t^b(\hat{y}) \\
= & ((e_j' y_t^b - \omega)/2) e_j' \dot{\Lambda}_t^b(\hat{y}) \\
= & ((e_j' y_t^b - \omega)/2) e_j' \big(b - \sum_{s=1}^{t} \mathbb{1}\{\Delta_s(\hat{y}) > 0\} a_s/t\big) \\
\geq & ((e_j' y_t^b - \omega)/2) e_j' \big(b - \sum_{s=1}^{t} \mathbb{1}\{u_s > a_s' e_j e_j' \hat{y}\} a_s/t\big) \\
\geq & ((e_j' y_t^b - \omega)/2) e_j' \big(b - \sum_{s=1}^{t} \mathbb{1}\{u > a' e_j \omega\} a_s/t\big) \\
= & ((e_j' y_t^b - \omega)/2) e_j' \dot{\Lambda}_t^b(\omega e_j).
\end{aligned}
$$

Since $e_j' y_t^b - \omega$ is positive, by assumption, it follows that $e_j' \dot{\Lambda}_t^b(\omega e_j)$ must be non-positive.

And now, I'll use this result to replace the shadow price with a simpler subgradient:

$$
\begin{aligned}
\Pr\big(\sup_{b \in B_\delta(\beta)} \|y_t^b\|^p > \gamma\big) &\leq \sum_{j=1}^{m} \Pr\big(\sup_{b \in B_\delta(\beta)} e_j' y_t^b > \gamma^{1/p}/\sqrt{m}\big) \\
&\leq \sum_{j=1}^{m} \Pr\big(\sup_{b \in B_\delta(\beta)} e_j' \dot{\Lambda}_t^b(e_j \gamma^{1/p}/\sqrt{m}) \leq 0\big).
\end{aligned}
$$

Next, we will bound the complex random variable in the last probability above with a simple binomial random variable. To that end, choose $\delta, \eta > 0$ so that $\eta \leq e_j' b$ for all $b \in B_\delta(\beta)$, in which case we have the following:

$$
\begin{aligned}
&\sup_{b \in B_\delta(\beta)} e_j' \dot{\Lambda}_t^b(e_j \gamma^{1/p}/\sqrt{m}) \\
&\qquad = \sup_{b \in B_\delta(\beta)} e_j' b - \sum_{s=1}^{t} \mathbb{1}\{u_s > a_s' e_j \gamma^{1/p}/\sqrt{m})\} e_j' a_s/t \\
&\qquad \geq \sup_{b \in B_\delta(\beta)} e_j' b - \sum_{s=1}^{t} \Big(\mathbb{1}\{e_j' a_s \leq e_j' b/2\}(e_j' b/2)/t + \mathbb{1}\{e_j' a_s > e_j' b/2\} \mathbb{1}\{u_s > a_s' e_j \gamma^{1/p}/\sqrt{m})\} \alpha/t\Big) \\
&\qquad \geq \sup_{b \in B_\delta(\beta)} e_j' b/2 - \sum_{s=1}^{t} \mathbb{1}\{u_s > e_j' b \gamma^{1/p}/(2\sqrt{m})\} \alpha/t \\
&\qquad \geq \eta/2 - \xi_t \alpha/t,
\end{aligned}
$$

where $\xi_t \equiv \sum_{s=1}^t \mathbb{1}\{u_s > \eta\gamma^{1/p}/(2\sqrt{m})\}$ is a binomial(t, ρ_γ), with $\rho_\gamma \equiv \Pr(u_1 > \eta\gamma^{1/p}/(2\sqrt{m}))$. Further, since $\mathrm{E}(u_1) \leq \infty$, we must have $\rho_\gamma \leq \gamma^{-1/p}$, for sufficiently large γ. Hence, combining the previous two results with the binomial Chernoff bound yields the following for sufficiently large γ:

$$\begin{aligned}
\Pr\big(\sup_{b\in B_\delta(\beta)} \|y_t^b\|^p > \gamma\big) &\leq \sum_{j=1}^m \Pr(\eta/2 - \xi_t\alpha/t \leq 0) \\
&= m\Pr(\xi_t \geq t\eta/(2\alpha)) \\
&\leq m\exp\Big(-\frac{t\eta}{2\alpha}\Big(\log\frac{\eta}{2\alpha\rho_\gamma} - 1\Big)\Big) \\
&\leq m\exp\Big(-\frac{t\eta}{4\alpha}\log\frac{\eta\gamma^{1/p}}{2\alpha}\Big) \\
&= m(\eta/(2\alpha))^{\frac{-t\eta}{4\alpha}}\gamma^{\frac{-t\eta}{4p\alpha}},
\end{aligned}$$

where the penultimate line supposes that γ is large enough to satisfy $\log(\frac{\eta\gamma^{1/p}}{2\alpha})/2 \geq 1$. Now choosing γ large enough to satisfy the previous result and large enough to ensure that $\|y_t^b\|^p \geq \gamma$ implies $y_t^b \notin B_\epsilon(y_\infty^b)$ yields the following:

$$\begin{aligned}
&\mathrm{E}(\sup_{b\in B_\delta(\beta)} \mathbb{1}\{y_t^b \notin B_{\gamma^{1/p}}(0)\}\|y_t^b\|^p) \\
&\leq \gamma\Pr\big(\sup_{b\in B_\delta(\beta)} \|y_t^b\|^p \geq \gamma\big) + \int_{x=\gamma}^\infty \Pr\big(\sup_{b\in B_\delta(\beta)} \|y_t^b\|^p > x\big)dx \\
&\leq \gamma\Pr\big(\sup_{b\in B_\delta(\beta)} \|y_t^b - y_\infty^b\| \geq \epsilon\big) + \int_{x=\gamma}^\infty m(\eta/(2\alpha))^{\frac{-t\eta}{4\alpha}}x^{\frac{-t\eta}{8\alpha}}dx \\
&\leq \gamma\Pr\big(\sup_{b\in B_\delta(\beta)} \|y_t^b - y_\infty^b\| \geq \epsilon\big) + \frac{m(\eta/(2\alpha))^{\frac{-t\eta}{4\alpha}}\gamma^{1-\frac{t\eta}{8\alpha}}}{\frac{t\eta}{8\alpha} - 1}.
\end{aligned}$$

The last expression above falls exponentially fast in t, by line (1), so this establishes line (2). □

Corollary 3 Proof. This is the $p = 0$ case of Proposition 6. □

Lemma 3 Proof. Consider an alternative martingale $\{\hat{b}_t\}_{t=n}^1$ in which $\hat{b}_n \equiv b_n$ and

$$\hat{b}_t \equiv \begin{cases} b_t & \hat{b}_{t+1} \in B_\delta(\beta), \\ \hat{b}_{t+1} & \hat{b}_{t+1} \notin B_\delta(\beta). \end{cases}$$

In other words, $\hat{b}_t$ tracks b_t until the first time that b_t departs $B_\delta(\beta)$, at which point $\hat{b}_t$ remains frozen in place. By design, $\hat{b}_t \in B_\delta(\beta)$ implies $b_t \in B_\delta(\beta)$, and hence $\Pr(b_t \notin B_\delta(\beta)) \leq \Pr(\hat{b}_t \notin B_\delta(\beta))$.

And, with this, the result follows from the Azuma–Hoeffding inequality, since $\|\hat{b}_t - \hat{b}_{t+1}\| \leq (\|\beta\| + \delta + \|\alpha\|)/t$:

$$
\begin{aligned}
&\Pr(b_t \notin B_\delta(\beta)) \\
&\quad \leq \sup_{N \geq t} \Pr(\hat{b}_t \notin B_\delta(\beta))) \\
&\quad \leq \sup_{N \geq t} \sum_{j=1}^{m} \Pr\left(|e_j'\hat{b}_t - e_j' b_n| \geq \delta/\sqrt{m}\right) \\
&\quad \leq \sup_{N \geq t} 2m \exp\Big(-\frac{\delta^2/m}{2\sum_{s=t}^{N-1}(\|\beta\| + \delta + \|\alpha\|)^2/s^2}\Big) \\
&\quad < 2m \exp\Big(-\frac{\delta^2}{2m(\|\beta\| + \delta + \|\alpha\|)^2 \int_{s=t-1}^{\infty} ds s^2}\Big) \\
&\quad < 2m \exp\Big(-\frac{\delta^2(t-1)}{2m(\|\beta\| + \delta + \|\alpha\|)^2}\Big). \qquad (4)
\end{aligned}
$$

□

Corollary 4 Proof. Let $\{\hat{b}_t\}_{t=n}^{1}$ be the alternative martingale defined in the proof of Lemma 3. Note that we have $\hat{b}_t \notin B_\delta(\beta)$ if and only if $t \leq \tau(\delta)$. And, with this, line (4) implies the result:

$$
\begin{aligned}
\mathrm{E}(\tau(\delta)) &= \sum_{s=1}^{n} \Pr(\tau(\delta) \geq s) \\
&= \sum_{t=1}^{n} \Pr(\hat{b}_t \notin B_\delta(\beta)) \\
&< \sum_{t=1}^{\infty} 2m \exp\Big(-\frac{\delta^2(t-1)}{2m(\|\beta\| + \delta + \|\alpha\|)^2}\Big) \\
&= O(1).
\end{aligned}
$$

□

Lemma 4 Proof. First, we can express the value function recursively:

$$
\bar{V}_t^b \equiv \begin{cases} \max_{x_t \in [0,1]} x_t u_t + \bar{V}_{t-1}^{\psi_t^b(x_t a_t)} & tb \geq a_t, \\ \bar{V}_{t-1}^{\psi_t^b(0)} & tb < a_t. \end{cases} \qquad (5)
$$

Second, since the shadow price weakly decreases with the inventory level, we have the following for

$x \in [0,1]$ and $tb \geq a_t$:

$$(1-x)a_t' y_{t-1}^{\psi_t^b(0)} \leq \bar{V}_{t-1}^{\psi_t^b(xa_t)} - \bar{V}_{t-1}^{\psi_t^b(a_t)} \leq (1-x)a_t' y_{t-1}^{\psi_t^b(a_t)} \tag{6}$$

$$\text{and} \quad x a_t' y_{t-1}^{\psi_t^b(0)} \leq \bar{V}_{t-1}^{\psi_t^b(0)} - \bar{V}_{t-1}^{\psi_t^b(xa_t)} \leq x a_t' y_{t-1}^{\psi_t^b(a_t)}. \tag{7}$$

Third, $\Delta_t(y_\infty^{b_t}) > 0$ and $tb_t \geq a_t$ imply $b_{t-1} = \psi_t^{b_t}(a_t)$ and hence $\hat{v}_{t-1} = \bar{V}_{t-1}^{\psi_t^{b_t}(a_t)} - \hat{R}_{t-1}$. Accordingly, lines (5)–(6) yield the following, when $\Delta_t(y_\infty^{b_t}) > 0$ and $tb_t \geq a_t$:

$$\begin{aligned}
\hat{R}_t =& \bar{V}_t^{b_t} - \hat{v}_t \\
=& \max_{x\in[0,1]} u_t x + \bar{V}_{t-1}^{\psi_t^{b_t}(xa_t)} - u_t - \hat{v}_{t-1} \\
=& \max_{x\in[0,1]} u_t(x-1) + \bar{V}_{t-1}^{\psi_t^{b_t}(xa_t)} - \bar{V}_{t-1}^{\psi_t^{b_t}(a_t)} + \hat{R}_{t-1} \\
\leq& \max_{x\in[0,1]} u_t(x-1) + (1-x)a_t' y_{t-1}^{\psi_t^{b_t}(a_t)} + \hat{R}_{t-1} \\
=& (a_t' y_{t-1}^{\psi_t^{b_t}(a_t)} - u_t)^+ + \hat{R}_{t-1} \\
=& \Delta_t(y_{t-1}^{\psi_t^{b_t}(a_t)})^- + \hat{R}_{t-1}.
\end{aligned} \tag{8}$$

Likewise, if $\Delta_t(y_\infty^{b_t}) \leq 0$ and $tb_t \geq a_t$ then (28), (5), and (7) yield

$$\hat{R}_t \leq \Delta_t(y_{t-1}^{\psi_t^{b_t}(0)})^+ + \hat{R}_{t-1}, \tag{9}$$

And if $tb_t < a_t$ then (28) and (5) yield

$$\hat{R}_t = \hat{R}_{t-1}. \tag{10}$$

Finally, since $\bar{V}_t^{b_t}$ can't exceed the sum of the remaining utilities, we must also have

$$\hat{R}_t \leq \sum_{s=1}^{t} u_s + \hat{R}_{t-1}. \tag{11}$$

Combining inequalities (8)–(11) inductively yields the result. □

Lemma 5 Proof. I will begin by bounding the expectation of the myopic regret's first term. Since

$\mathbb{1}\{b_t \notin B_{\delta/2}(\beta)\}$ is independent of $\sum_{s=1}^t u_s$, Lemma 3 indicates that there exists $C > 0$ for which

$$\begin{aligned}
&\mathrm{E}\big(\mathbb{1}\{b_t \notin B_{\delta/2}(\beta)\}\sum_{s=1}^t u_s\big) \\
&\quad = \Pr(b_t \notin B_{\delta/2}(\beta))\sum_{s=1}^t \mathrm{E}(u_s) \\
&\quad \le \exp(-Ct)t\,\mathrm{E}(u_1) \\
&\quad = o(1/t).
\end{aligned}$$

I will next bound the expectation of the myopic regret's second term. This second term is zero unless $b_t \in B_{\delta/2}(\beta)$. To streamline the math, I will henceforth suppress all $\mathbb{1}\{b_t \in B_{\delta/2}(\beta)\}$ indicator variables and implicitly suppose that the subsequent results condition on the event $b_t \in B_{\delta/2}(\beta)$.

To begin, note that $\|b_{t-1} - b_t\| \le \|2\beta + \alpha\|/(t-1)$ when $b_t \in B_{\delta/2}(\beta)$ and δ is sufficiently small. With this, Lemma 2 implies that we can choose δ small enough and t large enough to ensure that

$$\|y_\infty^{b_{t-1}} - y_\infty^{b_t}\| \le \|4\beta + 2\alpha\|/(\sigma_m^\beta(t-1)),$$

when $b_t \in B_{\delta/2}(\beta)$. Further, if t is sufficiently large then $b_t \in B_{\delta/2}(\beta)$ implies $b_{t-1} \in B_\delta(\beta)$, and hence

$$\|y_{t-1}^{b_{t-1}} - y_\infty^{b_{t-1}}\| \le \sup_{b \in B_\delta(\beta)} \|y_{t-1}^b - y_\infty^b\|.$$

Combining the previous two results yields the following, for sufficiently large t:

$$\begin{aligned}
\Delta_t(y_{t-1}^{b_{t-1}})^- &\le \Delta_t(y_\infty^{b_t} + \|y_\infty^{b_{t-1}} - y_\infty^{b_t}\|\iota + \|y_{t-1}^{b_{t-1}} - y_\infty^{b_{t-1}}\|\iota)^- \\
&\le \Delta_t(y_\infty^{b_t} + \xi_{t-1}\iota)^-, \qquad (12)\\
\text{where} \quad \xi_{t-1} &\equiv \|4\beta + 2\alpha\|/(\sigma_m^\beta(t-1)) + \sup_{b \in B_\delta(\beta)} \|y_{t-1}^b - y_\infty^b\|.
\end{aligned}$$

Note, it's easier to work with $\Delta_t(y_\infty^{b_t} + \xi_{t-1})^-$ than $\Delta_t(y_{t-1}^{b_{t-1}})^-$, because $y_\infty^{b_t} + \xi_{t-1}$ is independent of the random function Δ_t, whereas $y_{t-1}^{b_{t-1}}$ is not.

Now set δ small enough so that $b_t \in B_{\delta/2}(\beta)$ implies $y_\infty^{b_t} \in B_\epsilon(y_\infty^\beta)$. In this case, $y_\infty^{b_t} + \xi_{t-1}\iota \in$

$B_\epsilon(y_\infty^\beta)$ implies the following conditional expectation bound, by Line (12) and Lemma 4:

$$
\begin{aligned}
&\mathrm{E}\left(\mathbb{1}\{\Delta_t(y_\infty^{b_t}) > 0\}\Delta_t(y_{t-1}^{b_{t-1}})^- \mid b_t, \xi_{t-1}\right) \\
&\quad\leq \mathrm{E}\left(\mathbb{1}\{\Delta_t(y_\infty^{b_t}) > 0\}\Delta_t(y_\infty^{b_t} + \xi_{t-1}\iota)^- \mid b_t, \xi_{t-1}\right) \\
&\quad\leq 2\sigma_1^\beta \|\xi_{t-1}\iota\|^2 \\
&\quad= 2m\sigma_1^\beta \xi_{t-1}^2.
\end{aligned}
$$

Conversely, if $y_\infty^{b_t} + \xi_{t-1}\iota \notin B_\epsilon(y_\infty^\beta)$ then line (12) yields the following conditional expectation bound:

$$
\begin{aligned}
&\mathrm{E}\left(\mathbb{1}\{\Delta_t(y_\infty^{b_t}) > 0\}\Delta_t(y_{t-1}^{b_{t-1}})^- \mid b_t,\ \xi_{t-1}\right) \\
&\quad\leq \mathrm{E}\left(\Delta_t(y_\infty^{b_t} + \xi_{t-1}\iota)^- \mid b_t,\ \xi_{t-1}\right) \\
&\quad\leq \mathrm{E}\left(a_t'(y_\infty^{b_t} + \xi_{t-1}\iota) \mid b_t,\ \xi_{t-1}\right) \\
&\quad\leq \|\alpha\| \|y_\infty^{b_t} + \xi_{t-1}\iota\| \\
&\quad\leq \|\alpha\|(2\|y_\infty^\beta\| + \xi_{t-1}\sqrt{m}).
\end{aligned}
$$

For the final line, I suppose δ is small enough to ensure that $b_t \in B_{\delta/2}(\beta)$ implies $\|y_\infty^{b_t}\| \leq 2\|y_\infty^\beta\|$. Now, combining the previous two results yields the following, for sufficiently large t and small ϵ and δ:

$$
\begin{aligned}
&\mathrm{E}\left(\mathbb{1}\{\Delta_t(y_\infty^{b_t}) > 0\}\Delta_t(y_{t-1}^{b_{t-1}})^-\right) \\
&\quad\leq \mathrm{E}\left(\mathbb{1}\{y_\infty^{b_t} + \xi_{t-1}\iota \in B_\epsilon(y_\infty^\beta)\} 2m\sigma_1^\beta \xi_{t-1}^2\right) \\
&\qquad + \mathrm{E}\left(\mathbb{1}\{y_\infty^{b_t} + \xi_{t-1}\iota \notin B_\epsilon(y_\infty^\beta)\}\|\alpha\|(2\|y_\infty^\beta\| + \xi_{t-1}\sqrt{m})\right) \\
&\quad\leq 2m\sigma_1^\beta\,\mathrm{E}(\xi_{t-1}^2) + 2\|\alpha\|\|y_\infty^\beta\|\Pr(y_\infty^{b_t} + \xi_{t-1}\iota \notin B_\epsilon(y_\infty^\beta)) \\
&\qquad + \sqrt{m}\|\alpha\|\,\mathrm{E}\left(\mathbb{1}\{y_\infty^{b_t} + \xi_{t-1}\iota \notin B_\epsilon(y_\infty^\beta)\}\xi_{t-1}\right) \\
&\quad\leq 2m\sigma_1^\beta\,\mathrm{E}(\xi_{t-1}^2) + 2\|\alpha\|\|y_\infty^\beta\|\Pr(\xi_{t-1} \geq \epsilon/(2\sqrt{m})) \\
&\qquad + \sqrt{m}\|\alpha\|\,\mathrm{E}\left(\mathbb{1}\{\xi_{t-1} \geq \epsilon/(2\sqrt{m})\}\xi_{t-1}\right).
\end{aligned}
$$

The last line holds because we can make δ small enough so that $b_t \in B_{\delta/2}(\beta)$ implies $y_\infty^{b_t} \in B_{\epsilon/2}(y_\infty^\beta)$, in which case $b_t \in B_{\delta/2}(\beta)$ and $y_\infty^{b_t} + \xi_{t-1}\iota \notin B_\epsilon(y_\infty^\beta)$ imply that $\|\xi_{t-1}\iota\| \geq \epsilon/2$, and hence that $\xi_{t-1} \geq \epsilon/(2\sqrt{m})$. Finally, Proposition 4, Corollary 3, and Proposition 6 respectively establish that the first, second, and third terms of the last expression above are $O(1/t)$.

Finally, the same argument yields an analogous bound for the expectation of the myopic regret's third term. ∎

Lemma 6 Proof. Let $\{\hat{b}_t\}_{t=n}^1$ denote the inventory process defined in the proof of Lemma 3, but derived from from Algorithm 3's b_t values. Just to remind you, the $\{\hat{b}_t\}_{t=n}^1$ process tracks the $\{b_t\}_{t=n}^1$ process until time $\tau(\delta)$—i.e., until Algorithm 3's b_t values first depart $B_\delta(\beta)$—at which point the process freezes in place. The $\{\hat{b}_t\}_{t=n}^1$ process will be easier to study because a constant multiple of t bounds its innovations. And since $\hat{b}_t \in B_\delta(\beta)$ implies $b_t \in B_\delta(\beta)$, it will suffice to establish the concentration of measure for $\hat{b}_t$.

I will bound the distance between $\hat{b}_t$ and β with the following inequality:

$$\|\hat{b}_t - \beta\| \leq \|\hat{b}_{\tau(\delta/2)+1} - \beta\| + \|\xi_t\| + \sum_{s=t}^{\tau(\delta/2)} \|\operatorname{E}(\hat{b}_s \mid \hat{b}_{s+1}) - \hat{b}_{s+1}\|, \tag{13}$$
$$\text{where} \quad \xi_t \equiv \sum_{s=t}^{\tau(\delta/2)} \hat{b}_s - \operatorname{E}(\hat{b}_s \mid \hat{b}_{s+1}).$$

I cap the sums at time $\tau(\delta/2)$ to give our look-back shadow prices a sufficiently large sample. Indeed, a sample with $n - \tau(\delta/2)$ observations will comprise enough data to ensure that the look-back shadow prices—and hence the $\hat{b}_t$ values—are well-behaved. More specifically, I will show that $n - \tau(\delta/2) = \Theta(n)$ by showing that there exists $\gamma < 1$ that satisfies

$$\tau(\delta/2) + 1 \leq \gamma n. \tag{14}$$

To see this, note that period-t's resource vector satisfies

$$(n\beta - (n-t)\alpha)/t \leq \underbrace{\big(n\beta - \sum_{s=t}^{n} x_s a_s\big)/t}_{=b_t} \leq n\beta/t,$$

where the lower bound is within $\delta/2$ of β unless $t \leq \frac{n}{1+\delta/(2\|\alpha-\beta\|)}$, and the upper bound is within $\delta/2$ of β unless $t \leq \frac{n}{1+\delta/(2\|\beta\|)}$. Hence, if $\|\alpha - \beta\| \geq \|\beta\|$, which we can suppose without loss of generality, then $b_t \notin B_{\delta/2}(\beta)$ implies $t \leq \frac{n}{1+\delta/(2\|\alpha-\beta\|)}$.

I will now use (13) to inductively prove that there exists $C > 0$ such that

$$\Pr\big(\max_{s=t}^{n} \|\hat{b}_s - \beta\| > \delta\big) \leq (\tau(\delta/2) + 1 - t)\big(2\exp(-Ct) + 2n\exp(-C(1-\gamma)\sqrt{n})\big), \tag{15}$$

for all sufficiently large $t \leq n$. Initializing our induction will be simple: by definition, we have

$\Pr(\hat{b}_t \in B_\delta(\beta)) = 1$ for $t \geq \tau(\delta/2)+1$, which establishes the base case. However, establishing the inductive step will require unraveling the knotty relationship between look-back shadow prices and inventory vectors. Specifically, showing that $\|\hat{b}_t - \beta\|$ is small for $t \leq \tau(\delta/2)$ will require showing that $\|\operatorname{E}(\hat{b}_s \mid \hat{b}_{s+1}) - \hat{b}_{s+1}\|$ is small for all $s \in \{t, \cdots, \tau(\delta/2)\}$, which in turn will require showing that $\|\underleftarrow{y}_s^{\hat{b}_s} - y_\infty^{\hat{b}_s}\|$ is small for all $s \in \{t+1, \cdots, \tau(\delta/2)+1\}$, which in turn will require showing that $\|\hat{b}_s - \beta\|$ is small for all $s \in \{t+1, \cdots, \tau(\delta/2)+1\}$.

I will now that if (15) holds for sufficiently large $t+1 \leq \tau(\delta/2)+1$, then there a suitably high probability that

$$\begin{aligned} \|\hat{b}_{\tau(\delta/2)+1} - \beta\| &\leq \delta/2, \\ \|\xi_t\| &\leq \delta/4, \qquad\qquad (16) \\ \text{and} \quad \sum_{s=t}^{\tau(\delta/2)} \|\operatorname{E}(\hat{b}_s \mid \hat{b}_{s+1}) - \hat{b}_{s+1}\| &\leq \delta/4, \end{aligned}$$

which with line (13) will establish induction. Note that the first inequality in display (16) holds by the definition of $\tau(\delta/2)$, so we will only have to concern ourselves with the latter two inequalities.

I will now show that the second inequality in display (16) holds with high probability, conditional on $\hat{b}_s \in B_\delta(\beta)$ for all $s \in \{t+1, \cdots, \tau(\delta/2)+1\}$. Since $\{\xi_t\}_{t=\tau(\delta/2)}^1$ is a martingale that satisfies $\|\xi_t - \xi_{t+1}\| = \|\hat{b}_t - \operatorname{E}(\hat{b}_t \mid \hat{b}_{t+1})\| \leq (\|\beta\| + \delta + \|\alpha\|)/t$, by design, the argument underlying line (4) analogously implies that there exists $C > 0$ such that $\Pr(\|\xi_t\| > \delta/4) \leq \exp(-Ct)$, for all sufficiently large t. And since $\Pr(A|B) = \Pr(A \cap B)/\Pr(B) \leq \Pr(A)/\Pr(B)$, it follows that

$$\begin{aligned} &\Pr\left(\|\xi_t\| > \delta/4 \mid \max_{s=t+1}^{\tau(\delta/2)+1} \|\hat{b}_s - \beta\| \leq \delta\right) \\ &\leq \frac{\Pr(\|\xi_t\| > \delta/4)}{\Pr\left(\max_{s=t+1}^{\tau(\delta/2)+1} \|\hat{b}_s - \beta\| \leq \delta\right)} \\ &\leq 2\exp(-Ct). \qquad\qquad (17) \end{aligned}$$

Note, the last line holds because $\Pr(\max_{s=t+1}^{\tau(\delta/2)+1} \|\hat{b}_s - \beta\| \leq \delta) \geq 1/2$, by our inductive hypothesis.

I will now show that the third inequality in display (16) holds with high probability, conditional on $\hat{b}_s \in B_\delta(\beta)$ for all $s \in \{t+1, \cdots, \tau(\delta/2)+1\}$. This step will take more work. First note that

$\hat{b}_{s+1} \in B_\delta(\beta)$ implies $\hat{b}_{s+1} = b_{s+1}$ and $\hat{b}_s = b_s$, and thus implies

$$\begin{aligned}
&\mathrm{E}(\hat{b}_s \mid \hat{b}_{s+1}) - \hat{b}_{s+1} \\
&\quad = \mathrm{E}(b_s \mid b_{s+1}) - b_{s+1} \\
&\quad = ((s+1)b_{s+1} - \mathrm{E}(\mathbb{1}\{\Delta_{s+1}(\underleftarrow{y}_{s+1}^{b_{s+1}}) > 0\}a_{s+1} \mid b_{s+1}))/s - b_{s+1} \\
&\quad = ((s+1)b_{s+1} - b_{s+1} + \dot{\Lambda}_\infty^{b_{s+1}}(\underleftarrow{y}_{s+1}^{b_{s+1}}))/s - b_{s+1} \\
&\quad = \dot{\Lambda}_\infty^{b_{s+1}}(\underleftarrow{y}_{s+1}^{b_{s+1}})/s \\
&\quad = \ddot{\Lambda}_\infty(y_\infty^{b_{s+1}})(\underleftarrow{y}_{s+1}^{b_{s+1}} - y_\infty^{b_{s+1}})/s + o(\|\underleftarrow{y}_{s+1}^{b_{s+1}} - y_\infty^{b_{s+1}}\|)/s \\
&\quad = \ddot{\Lambda}_\infty(y_\infty^{\hat{b}_{s+1}})(\underleftarrow{y}_{s+1}^{\hat{b}_{s+1}} - y_\infty^{\hat{b}_{s+1}})/s + o(\|\underleftarrow{y}_{s+1}^{\hat{b}_{s+1}} - y_\infty^{\hat{b}_{s+1}}\|)/s,
\end{aligned}$$

where the penultimate line holds by Lemma 1, since $\dot{\Lambda}_\infty^{b_{s+1}}(y_\infty^{b_{s+1}}) = 0$ when $b_{s+1} \in B_\delta(\beta)$ and δ is small. Thus, we can choose n sufficiently small so that $\max_{s=t+1}^{\tau(\delta/2)+1} \|\hat{b}_s - \beta\| \leq \delta$ and $\max_{s=t+1}^{\tau(\delta/2)+1} \|\underleftarrow{y}_s^{\hat{b}_s} - y_\infty^{\hat{b}_s}\| \leq n^{-1/4}$ imply

$$\begin{aligned}
&\sum_{s=t}^{\tau(\delta/2)} \|\mathrm{E}(\hat{b}_s \mid \hat{b}_{s+1}) - \hat{b}_{s+1}\| \\
&\quad \leq \sum_{s=t}^{\tau(\delta/2)} \|\ddot{\Lambda}_\infty(y_\infty^{\hat{b}_{s+1}})(\underleftarrow{y}_{s+1}^{\hat{b}_{s+1}} - y_\infty^{\hat{b}_{s+1}})\|/s + o(\|\underleftarrow{y}_{s+1}^{\hat{b}_{s+1}} - y_\infty^{\hat{b}_{s+1}}\|)/s \\
&\quad \leq \sum_{s=t}^{\tau(\delta/2)} 2\sigma_1^\beta \|\underleftarrow{y}_{s+1}^{\hat{b}_{s+1}} - y_\infty^{\hat{b}_{s+1}}\|/s + o(\|\underleftarrow{y}_{s+1}^{\hat{b}_{s+1}} - y_\infty^{\hat{b}_{s+1}}\|)/s \\
&\quad \leq \sum_{s=t}^{\tau(\delta/2)} 3\sigma_1^\beta n^{-1/4}/s \\
&\quad \leq \delta/4. \qquad (18)
\end{aligned}$$

Note, the third line holds because the largest singular value of $\ddot{\Lambda}_\infty(y_\infty^{\hat{b}_{s+1}})$ is less than twice the largest singular value of $\ddot{\Lambda}_\infty(y_\infty^\beta)$ when $\hat{b}_{s+1} \in B_\delta(\beta)$ and δ is small, and the fourth line holds because the little-o term is less than $\sigma_1^\beta n^{-1/4}$ when $\|\underleftarrow{y}_{s+1}^{\hat{b}_{s+1}} - y_\infty^{\hat{b}_{s+1}}\| \leq n^{-1/4}$ and n is large.

Further, we can use Corollary 3 upper bound the probability that $\max_{s=t+1}^{\tau(\delta/2)+1} \|\underleftarrow{y}_s^{\hat{b}_s} - y_\infty^{\hat{b}_s}\| > n^{-1/4}$, conditional on $\max_{s=t+1}^{\tau(\delta/2)+1} \|\hat{b}_s - \beta\| \leq \delta$. Specifically, combining this corollary with line (14) and our inductive hypothesis yields the following, for some $C > 0$ and all sufficiently large n:

$$
\begin{aligned}
\Pr\Big(& \max_{s=t+1}^{\tau(\delta/2)+1} \| \underleftarrow{y}_s^{\hat{b}_s} - y_\infty^{\hat{b}_s} \| > n^{-1/4} \mid \max_{s=t+1}^{\tau(\delta/2)+1} \|\hat{b}_s - \beta\| \le \delta \Big) \\
&\le \sum_{s=t+1}^{\tau(\delta/2)+1} \Pr\Big(\sup_{b \in B_\delta(\beta)} \| \underleftarrow{y}_s^{b} - y_\infty^{b} \| > n^{-1/4} \mid \max_{s=t+1}^{\tau(\delta/2)+1} \|\hat{b}_s - \beta\| \le \delta \Big) \\
&\le \sum_{s=t+1}^{\tau(\delta/2)+1} \frac{\Pr\big(\sup_{b \in B_\delta(\beta)} \| \underleftarrow{y}_s^{b} - y_\infty^{b} \| > n^{-1/4} \big)}{\Pr\big(\max_{s=t+1}^{\tau(\delta/2)} \|\hat{b}_s - \beta\| \le \delta \big)} \\
&\le \sum_{s=t+1}^{\tau(\delta/2)+1} \frac{\exp(-Cn^{-1/2}(n-s))}{1/2} \\
&\le 2n \exp(-Cn^{-1/2}(n - \tau(\delta/2) - 1)) \\
&\le 2n \exp(-C(1-\gamma)\sqrt{n}).
\end{aligned} \tag{19}
$$

And now, finally, we can combine lines (13), (17), (18), and (19) to establish that

$$
\begin{aligned}
\Pr(& \hat{b}_t \notin B_\delta(\beta) \mid \max_{s=t+1}^{n} \|\hat{b}_s - \beta\| \le \delta) \\
&\le \Pr\big(\|\xi_t\| > \delta/4 \mid \max_{s=t+1}^{n} \|\hat{b}_s - \beta\| \le \delta \big) \\
&\qquad + \Pr\Big(\sum_{s=t}^{\tau(\delta/2)} \| \mathrm{E}(\hat{b}_s \mid \hat{b}_{s+1}) - \hat{b}_{s+1} \| > \delta/4 \mid \max_{s=t+1}^{n} \|\hat{b}_s - \beta\| \le \delta \Big) \\
&\le 2\exp(-Ct) + \Pr\Big(\max_{s=t+1}^{\tau(\delta/2)+1} \| \underleftarrow{y}_s^{\hat{b}_s} - y_\infty^{\hat{b}_s} \| > n^{-1/4} \mid \max_{s=t+1}^{n} \|\hat{b}_s - \beta\| \le \delta \Big) \\
&\le 2\exp(-Ct) + 2n\exp(-C(1-\gamma)\sqrt{n}).
\end{aligned}
$$

And with our inductive hypothesis, this implies that

$$
\begin{aligned}
\Pr\big(& \max_{s=t}^{n} \|\hat{b}_s - \beta\| > \delta \big) \\
&= \Pr\big(\max_{s=t+1}^{n} \|\hat{b}_s - \beta\| > \delta \big) + \Pr(\hat{b}_t \notin B_\delta(\beta) \mid \max_{s=t+1}^{n} \|\hat{b}_s - \beta\| \le \delta) \\
&\le (\tau(\delta/2) + 1 - t - 1)(2\exp(-C(t+1)) + 2n\exp(-C(1-\gamma)\sqrt{n})) \\
&\qquad + 2\exp(-Ct) + 2n\exp(-C(1-\gamma)\sqrt{n}) \\
&\le (\tau(\delta/2) + 1 - t)\big(2\exp(-Ct) + 2n\exp(-C(1-\gamma)\sqrt{n})\big).
\end{aligned}
$$

$\square$

Lemma 7 Proof. This result follows from the argument used to establish Lemma 4. □

Lemma 8 Proof. This proof will closely follow the proof of Lemma 5. For example, $\mathrm{E}(\mathbb{1}\{b_t \notin B_{\delta/2}(\beta)\} \sum_{s=1}^{t} u_s) = o(1/t)$ trivially follows from Lemma 6, as it previously followed from Lemma 3.

Bounding the expectation of the myopic regret's second term will prove more difficult. As in the proof of Lemma 5, I will henceforth suppress all $\mathbb{1}\{b_t \in B_{\delta/2}(\beta)\}$ indicator variables and implicitly suppose that the subsequent results condition on the event $b_t \in B_{\delta/2}(\beta)$.

Let ξ_{t-1} be as defined in the proof of Lemma 5, and analogously define $\underleftarrow{\xi}_{t+1} \equiv \sup_{b \in B_{\delta/2}(\beta)} \| \underleftarrow{y}_t^{\,b} - y_\infty^b \|$. I subscript this latter variable with $t+1$ because $\underleftarrow{y}_t^{\,b}$ is determined by that time. Note, that $b_t \in B_{\delta/2}(\beta)$ implies $\underleftarrow{y}_t^{\,b_t} \geq y_\infty^{b_t} - \underleftarrow{\xi}_{t+1}\iota$, and hence that $\Delta_t(\underleftarrow{y}_t^{\,b_t}) \leq \Delta_t(y_\infty^{b_t} - \underleftarrow{\xi}_{t+1}\iota)$. With this, Line (12) implies the following:

$$
\begin{aligned}
&\mathrm{E}\left(\mathbb{1}\{\Delta_t(\underleftarrow{y}_t^{\,b_t}) > 0\}\Delta_t(y_{t-1}^{b_{t-1}})^-\right) \\
&\quad \leq \mathrm{E}\left(\mathbb{1}\{\Delta_t(y_\infty^{b_t} - \underleftarrow{\xi}_{t+1}\iota) > 0\}\Delta_t(y_\infty^{b_t} + \xi_{t-1}\iota)^-\right).
\end{aligned} \tag{20}
$$

Now suppose $y_\infty^{b_t} + \xi_{t-1}\iota \in B_\epsilon(y_\infty^\beta)$ and $y_\infty^{b_t} - \underleftarrow{\xi}_t\iota \in B_\epsilon(y_\infty^\beta)$. In this case, Line (20) and Lemma 4 yield the following conditional expectation bound:

$$
\begin{aligned}
&\mathrm{E}\left(\mathbb{1}\{\Delta_t(\underleftarrow{y}_t^{\,b_t}) > 0\}\Delta_t(y_{t-1}^{b_{t-1}})^- \mid b_t, \xi_{t-1}, \underleftarrow{\xi}_{t+1}\right) \\
&\quad \leq \mathrm{E}\left(\mathbb{1}\{\Delta_t(y_\infty^{b_t} - \underleftarrow{\xi}_{t+1}\iota) > 0\}\Delta_t(y_\infty^{b_t} + \xi_{t-1}\iota)^- \mid b_t, \xi_{t-1}, \underleftarrow{\xi}_{t+1}\right) \\
&\quad \leq 2\sigma_1^\beta \|\underleftarrow{\xi}_{t+1}\iota - \xi_{t-1}\iota\|^2 \\
&\quad \leq 4m\sigma_1^\beta(\underleftarrow{\xi}_{t+1}^{\,2} + \xi_{t-1}^2).
\end{aligned}
$$

Conversely, if $y_\infty^{b_t} + \xi_{t-1}\iota \notin B_\epsilon(y_\infty^\beta)$ or $y_\infty^{b_t} - \underleftarrow{\xi}_t\iota \notin B_\epsilon(y_\infty^\beta)$ then Line (20) yields the following conditional expectation bound:

$$
\begin{aligned}
&\mathrm{E}\left(\mathbb{1}\{\Delta_t(\underleftarrow{y}_t^{\,b_t}) > 0\}\Delta_t(y_{t-1}^{b_{t-1}})^- \mid b_t, \xi_{t-1}, \underleftarrow{\xi}_{t+1}\right) \\
&\quad \leq \mathrm{E}\left(\Delta_t(y_\infty^{b_t} + \xi_{t-1}\iota)^- \mid b_t, \xi_{t-1}, \underleftarrow{\xi}_{t+1}\right) \\
&\quad \leq \mathrm{E}\left(a_t'(y_\infty^{b_t} + \xi_{t-1}\iota) \mid b_t, \xi_{t-1}\right) \\
&\quad \leq \|\alpha\|\|y_\infty^{b_t} + \xi_{t-1}\iota\| \\
&\quad \leq \|\alpha\|(2\|y_\infty^\beta\| + \xi_{t-1}m).
\end{aligned}
$$

For the last line, I suppose δ is small enough to ensure that $b_t \in B_{\delta/2}(\beta)$ implies $\|y_\infty^{b_t}\| \leq 2\|y_\infty^\beta\|$.

Now, define $\mathcal{E} \equiv \mathbb{1}\{y_\infty^{b_t} + \xi_{t-1}\iota \in B_\epsilon(y_\infty^\beta)\}\mathbb{1}\{y_\infty^{b_t} - \underleftarrow{\xi}_t\iota \in B_\epsilon(y_\infty^\beta)\}$. With this, the previous two results yield the following, for sufficiently large t and small ϵ and δ:

$$
\begin{aligned}
&\mathrm{E}\left(\mathbb{1}\{\Delta_t(\underleftarrow{y}_t^{b_t}) > 0\}\Delta_t(y_{t-1}^{b_{t-1}})^-\right) \\
&\quad \leq \mathrm{E}\left(\mathcal{E}4m\sigma_1^\beta(\underleftarrow{\xi}_{t+1}^2 + \xi_{t-1}^2)\right) + \mathrm{E}\left((1-\mathcal{E})\|\alpha\|(2\|y_\infty^\beta\| + \xi_{t-1}m)\right) \\
&\quad \leq 4m\sigma_1^\beta\,\mathrm{E}(\underleftarrow{\xi}_{t+1}^2 + \xi_{t-1}^2) \\
&\qquad + 2\|\alpha\|\|y_\infty^\beta\|\Pr(y_\infty^{b_t} + \xi_{t-1}\iota \notin B_\epsilon(y_\infty^\beta)) \\
&\qquad + 2\|\alpha\|\|y_\infty^\beta\|\Pr(y_\infty^{b_t} - \underleftarrow{\xi}_t\iota \notin B_\epsilon(y_\infty^\beta)) \\
&\qquad + m\,\mathrm{E}\left(\mathbb{1}\{y_\infty^{b_t} + \xi_{t-1}\iota \notin B_\epsilon(y_\infty^\beta)\}\xi_{t-1}\right) \\
&\qquad + m\Pr\left(y_\infty^{b_t} - \underleftarrow{\xi}_t\iota \notin B_\epsilon(y_\infty^\beta)\right)\mathrm{E}(\xi_{t-1}).
\end{aligned}
$$

Note, I can separate $\underleftarrow{\xi}_t$ and ξ_{t-1} in the last term of the final expression because these variables are independent of one another. And each of the terms in this final expression is either $O(1/t)$ or $O(1/(n-t))$, by the argument used at the end of Lemma 5.

Finally, the argument above yields an analogous bound for the expectation of the myopic regret's third term. □

Corollary 5 Proof. Copy the proof of Corollary 4. □

Lemma 9 Proof. I will begin with a high-level plan of attack. The main idea of the proof is that allocating $t(\beta+\xi)$ units of inventory to the first t periods leaves us with only $(n-t)(\beta - \frac{t}{n-t}\xi)$ units for the last $n-t$ periods, so $b_t = \beta + \xi$ must imply

$$
\begin{aligned}
v_n^\beta &\leq \bar{V}_t^{\beta+\xi} + \underleftarrow{\bar{V}}_t^{\beta - \frac{t}{n-t}\xi} \\
&= t\Lambda_t^{\beta+\xi}(y_t^{\beta+\xi}) + (n-t)\underleftarrow{\Lambda}_t^{\beta - \frac{t}{n-t}\xi}(\underleftarrow{y}_t^{\beta - \frac{t}{n-t}\xi}) \\
&\equiv \bar{\Lambda}_t^\xi
\end{aligned}
\tag{21}
$$

where $\underleftarrow{\bar{V}}_t^b$, $\underleftarrow{\Lambda}_t^b$, and $\underleftarrow{y}_t^b$ and are equivalent to $\bar{V}_{n-t}^b$, Λ_{n-t}^b, and y_{n-t}^b, but with the order of the customers reversed. Accordingly, it follows that

$$
\hat{R}_n \geq \bar{V}_n^\beta - \bar{\Lambda}_t^{b_t - \beta} = n\Lambda_n^\beta(y_n^\beta) - \bar{\Lambda}_t^{b_t-\beta}. \tag{22}
$$

And the expression on the right should be large when b_t meaningfully deviates from β since, in the limit, we have

$$n\Lambda_\infty^\beta(y_\infty^\beta) - \bar{\Lambda}_\infty^\xi \geq Ct\min(\|\xi\|^2, 1), \tag{23}$$

for some $C > 0$, where[1]

$$\bar{\Lambda}_\infty^\xi \equiv t\Lambda_\infty^{\beta+\xi}(y_\infty^{\beta+\xi}) + (n-t)\Lambda_\infty^{\beta-\frac{t}{n-t}\xi}(y_\infty^{\beta-\frac{t}{n-t}\xi}). \tag{24}$$

Line (23) suggests that reserving $t(\beta+\xi)$ units of inventory for the first t periods should sacrifice $\Omega(t)$ units of value when ξ non-negligible. I will leverage this fact to show that $\hat{R}_n$ is almost always large when $\|b_t - \beta\|$ is non-negligible. And since $\mathrm{E}(\hat{R}_n)$ is relatively small, by Theorem 1, it follows that $\|b_t - \beta\|$ is usually negligible.

Before delving into the details, I will provide a more thorough proof sketch. The proof will have five steps. The first derives limiting bound (23), our only tool for establishing the cost of b_t diverging from β. Now with (23), it's relatively easy to lower bound the regret when $b_t = \beta + \xi$, for some specific ξ that lies outside of a ball of the origin. But that's not enough, as we must lower bound this regret when $b_t = \beta + \xi$, for *any* ξ that lies outside of a ball of the origin. To create such a uniform result, the second part of the proof bounds $\bar{\Lambda}_t^\xi$ for all $\xi \in \mathbb{R}^m$ in terms of $\bar{\Lambda}_t^\zeta$, $y_t^{\beta+\zeta}$, and $\underleftarrow{y}_t^{\beta-\frac{t}{n-t}\zeta}$, for some given given $\zeta \in \mathbb{R}^m$, and the third part uses this bound to show that $\sup_{\xi\notin B_{2\sqrt{m}\epsilon}(0)} \bar{\Lambda}_t^\xi$ is usually smaller than $\max_{k\in\{-1,1\}}\max_{j\in[m]} \bar{\Lambda}_t^{k\epsilon\omega_j^\beta}$, where ϵ is a small positive number and $\{\omega_i^\beta\}_{i\in[m]}$ are the orthonormal eigenvectors of $\ddot{\Lambda}_\infty(y_\infty^\beta)^{-1}$. Hence, the second and third steps of the proof collapse the relevant domain of $b_t - \beta$ from the infinite set $\mathbb{R}^m \setminus B_{2\sqrt{m}\epsilon}(0)$ to the finite set $\{k\epsilon\omega_j^\beta \mid k \in \{-1,1\} \times j \in [m]\}$. The fourth step of the proof uses our shadow price convergence results to show that $n\Lambda_n^\beta(y_n^\beta) - \max_{k\in\{-1,1\}}\max_{j\in[m]} \bar{\Lambda}_t^{k\epsilon\omega_j^\beta}$ is usually very large, and the last step combines this with the previous results to establish that $\xi = b_t - \beta$ must rarely fall outside of $B_{2\sqrt{m}\epsilon}(0)$.

To begin the proof, note that Lemmas 1 and 2 imply that $\Lambda_\infty^b(y_\infty^b)$ is concave in b, since $\frac{\partial^2}{\partial b^2}\Lambda_\infty^b(y_\infty^b) = \frac{\partial}{\partial b}y_\infty^b = -\ddot{\Lambda}_\infty(y_\infty^b)^{-1}$ is negative definite. This concavity implies that we can restrict attention to small ξ vectors since $\bar{\Lambda}_\infty^\xi$ decreases in the magnitude of ξ. But, more importantly, the concavity implies line (23), as I will now show.

Let $i \in [m]$ denote the index of the largest element of ξ, so that either $e_i'\xi = \|\xi\|_\infty$ or $-e_i'\xi = \|\xi\|_\infty$. Since the minus sign doesn't meaningfully affect the analysis, I will henceforth suppose

$e_i'\xi = \|\xi\|_\infty \equiv \gamma$, in which case

$$\bar{\Lambda}^{\xi}_{\infty} \leq \sup_{\{\zeta \in \mathbb{R}^m \mid e_i'\zeta = \gamma\}} \bar{\Lambda}^{\zeta}_{\infty}. \tag{25}$$

The solution to this optimization problem satisfies the following first-order conditions for some Lagrange multiplier λ:

$$\begin{aligned} 0 =& \tfrac{\partial}{\partial \zeta}\Big(t\Lambda^{\beta+\zeta}_{\infty}(y^{\beta+\zeta}_{\infty}) + (n-t)\Lambda^{\beta-\frac{t}{n-t}\zeta}_{\infty}(y^{\beta-\frac{t}{n-t}\zeta}_{\infty}) - \lambda(e_i'\zeta - \gamma)\Big) \\ =& t(y^{\beta+\zeta}_{\infty} - y^{\beta-\frac{t}{n-t}\zeta}_{\infty}) - \lambda e_i. \end{aligned}$$

Now we'll use Lemma 2 to differentiate this with respect to γ:

$$\begin{aligned} 0 &= \tfrac{\partial}{\partial \gamma} 0\big|_{\gamma=0} \\ &= \tfrac{\partial}{\partial \gamma}\Big(t(y^{\beta+\zeta}_{\infty} - y^{\beta-\frac{t}{n-t}\zeta}_{\infty}) - \lambda e_i\Big)\big|_{\gamma=0} \\ &= -\Big(t\ddot{\Lambda}_{\infty}(y^{\beta+\zeta}_{\infty})^{-1} + \frac{t^2}{n-t}\ddot{\Lambda}_{\infty}(y^{\beta-\frac{t}{n-t}\zeta}_{\infty})^{-1}\Big)\tfrac{\partial}{\partial \gamma}\zeta - e_i \tfrac{\partial}{\partial \gamma}\lambda\big|_{\gamma=0} \\ &= -\frac{nt}{n-t}\ddot{\Lambda}_{\infty}(y^{\beta}_{\infty})^{-1}\tfrac{\partial}{\partial \gamma}\zeta - e_i\tfrac{\partial}{\partial \gamma}\lambda\big|_{\gamma=0}. \end{aligned}$$

The last line holds because $\gamma = 0$ implies $\zeta = 0$. Combining the $e_i'\zeta = \gamma$ constraint with the expression above yields

$$-\frac{n-t}{nt}e_i'\ddot{\Lambda}_{\infty}(y^{\beta}_{\infty})e_i\tfrac{\partial}{\partial \gamma}\lambda\big|_{\gamma=0} = \tfrac{\partial}{\partial \gamma}e_i'\zeta\big|_{\gamma=0} = \tfrac{\partial}{\partial \gamma}\gamma\big|_{\gamma=0} = 1,$$

which implies that

$$\tfrac{\partial}{\partial \gamma}\lambda\big|_{\gamma=0} = \frac{-nt}{(n-t)e_i'\ddot{\Lambda}_{\infty}(y^{\beta}_{\infty})e_i}.$$

Further, since $\lambda = 0$ when $\gamma = 0$, by the concavity of $\Lambda^b_{\infty}(y^b_{\infty})$ in b, it follows that for sufficiently small γ we have

$$\lambda \leq \frac{-nt\gamma}{2(n-t)e_i'\ddot{\Lambda}_{\infty}(y^{\beta}_{\infty})e_i}.$$

By definition, our Lagrange multiplier also satisfies $\frac{\partial}{\partial \gamma}\sup_{\{\zeta \in \mathbb{R}^m \mid e_i'\zeta = \gamma\}} \bar{\Lambda}^{\zeta}_{\infty} = \lambda$, which with the result above yields the following, for sufficiently small γ:

$$
\begin{aligned}
n\Lambda_\infty^\beta(y_\infty^\beta) - &\sup_{\{\zeta\in\mathbb{R}^m \mid e_i'\zeta=\gamma\}} \bar\Lambda_\infty^\zeta \\
&= -\Big(\sup_{\{\zeta\in\mathbb{R}^m \mid e_i'\zeta=\gamma\}} \bar\Lambda_\infty^\zeta - \sup_{\{\zeta\in\mathbb{R}^m \mid e_i'\zeta=0\}} \bar\Lambda_\infty^\zeta\Big) \\
&= -\int_{g=0}^{\gamma} \frac{\partial}{\partial g} \sup_{\{\zeta\in\mathbb{R}^m \mid e_i'\zeta=g\}} \bar\Lambda_\infty^\zeta dg \\
&\geq -\int_{g=0}^{\gamma} \frac{-ntg}{2(n-t)e_i'\ddot\Lambda_\infty(y_\infty^\beta)e_i} dg \\
&= \frac{nt\gamma^2}{4(n-t)e_i'\ddot\Lambda_\infty(y_\infty^\beta)e_i} \\
&\geq \frac{t(\|\xi\|/\sqrt{m})^2}{4\max_{j\in[m]} e_j'\ddot\Lambda_\infty(y_\infty^\beta)e_j}.
\end{aligned}
$$

Note, the first line above holds because the concavity of $\Lambda_\infty^b(y_\infty^b)$ in b implies that $n\Lambda_\infty^\beta(y_\infty^\beta) = \sup_{\{\zeta\in\mathbb{R}^m \mid e_i'\zeta=0\}} \bar\Lambda_\infty^\zeta$, and the last line holds because $\gamma = \|\xi\|_\infty \geq \|\xi\|/\sqrt{m}$. Finally, combining the result above with line (25) yields line (23).

Second, I will now bound the difference between $\bar\Lambda_t^\xi$ and $\bar\Lambda_t^\zeta$ in terms of $y_t^{\beta+\zeta}$ and $\underleftarrow{y}_t^{\beta-\frac{t}{n-t}\zeta}$, which will enable us to invoke our shadow price convergence results. To this end, first note that

$$
\begin{aligned}
\bar\Lambda_t^\xi = &\max_{x\in[0,1]^n} \sum_{s=1}^{n} x_s u_s \\
\text{s. t.} \quad &\sum_{s=1}^{t} x_s a_s \leq t(\beta+\xi), \\
&\sum_{s=t+1}^{n} x_s a_s \leq (n-t)\big(\beta - \frac{t}{n-t}\xi\big).
\end{aligned}
$$

Since this linear program is concave in its constraints, $\bar\Lambda_t^\xi$ must be concave in ξ. Accordingly, the $\bar\Lambda_t^\xi$ function lies below the hyperplane characterized by supergradient $\frac{\partial}{\partial\zeta}\bar\Lambda_t^\zeta \equiv t y_t^{\beta+\zeta}\frac{\partial}{\partial\zeta}(\beta+\zeta) + (n-t)\underleftarrow{y}_t^{\beta-\frac{t}{n-t}\zeta}\frac{\partial}{\partial\zeta}\big(\beta-\frac{t}{n-t}\zeta\big) = t(y_t^{\beta+\zeta} - \underleftarrow{y}_t^{\beta-\frac{t}{n-t}\zeta})$:

$$
\bar\Lambda_t^\xi - \bar\Lambda_t^\zeta \leq t(\xi-\zeta)'\big(y_t^{\beta+\zeta} - \underleftarrow{y}_t^{\beta-\frac{t}{n-t}\zeta}\big). \tag{26}
$$

Third, I will use the preceding inequality to establish that

$$\Pr\big(\sup_{\xi \notin B_{2\sqrt{m}\epsilon}(0)} \bar{\Lambda}_t^{\xi} \le \max_{k\in\{-1,1\}} \max_{j\in[m]} \bar{\Lambda}_t^{k\epsilon\omega_j^{\beta}}\big) \ge 3/4, \tag{27}$$

for all sufficiently small $\epsilon > 0$ and large t. This bound is crucial, as it enables us to replace the infinite continuum of $\bar{\Lambda}_t^{\xi}$ values for all $\xi \notin B_{2\sqrt{m}\epsilon}(0)$, with the largest of the $2m$ values of $\bar{\Lambda}_t^{k\epsilon\omega_j^{\beta}}$. To begin, note that Lemma 2 yields the following, for $k \in \{-1, 1\}$ and small $\epsilon > 0$:

$$y_{\infty}^{\beta+k\epsilon\omega_i} - y_{\infty}^{\beta} = -k\epsilon\ddot{\Lambda}_{\infty}(y_{\infty}^{\beta})^{-1}\omega_i^{\beta} + o(\epsilon) = -k\epsilon\omega_i^{\beta}/\sigma_i^{\beta} + o(\epsilon).$$

Combining this with (26) yields the following, for $t \le n/2$:

$$\begin{aligned}
\bar{\Lambda}_t^{\xi} - \bar{\Lambda}_t^{k\epsilon\omega_i^{\beta}} &\le t(\xi - k\epsilon\omega_i^{\beta})'\big(y_t^{\beta+k\epsilon\omega_i^{\beta}} - \underleftarrow{y}_t^{\beta-\frac{t}{n-t}k\epsilon\omega_i^{\beta}}\big) \\
&= t(\xi - k\epsilon\omega_i^{\beta})'\big(y_{\infty}^{\beta+k\epsilon\omega_i^{\beta}} - y_{\infty}^{\beta} - y_{\infty}^{\beta-\frac{t}{n-t}k\epsilon\omega_i^{\beta}} + y_{\infty}^{\beta}\big) \\
&\quad + t(\xi - k\epsilon\omega_i^{\beta})'(y_t^{\beta+k\epsilon\omega_i^{\beta}} - y_{\infty}^{\beta+k\epsilon\omega_i^{\beta}}) - t(\xi - k\epsilon\omega_i^{\beta})'(\underleftarrow{y}_t^{\beta-\frac{t}{n-t}k\epsilon\omega_i} - y_{\infty}^{\beta-\frac{t}{n-t}k\epsilon\omega_i}) \\
&= \frac{-ntk\epsilon}{n-t}(\xi - k\epsilon\omega_i^{\beta})'\omega_i^{\beta}/\sigma_i^{\beta} + t\|\xi - k\epsilon\omega_i^{\beta}\|(o(\epsilon) + O_p(t^{-1/2})),
\end{aligned} \tag{28}$$

where the last line holds because $\|y_t^{\beta+\epsilon\omega_i^{\beta}} - y_{\infty}^{\beta+\epsilon\omega_i^{\beta}}\|$ and $\|\underleftarrow{y}_t^{\beta-\frac{t}{n-t}\epsilon\omega_i} - y_{\infty}^{\beta-\frac{t}{n-t}\epsilon\omega_i}\|$ are $O_p(t^{-1/2})$ when $t \le n/2$, by Proposition 4.

Now, to derive (27) from (28), let $\gamma_i \equiv \xi'\omega_i^{\beta}$, so that $\xi = \sum_{i=1}^m \gamma_i\omega_i^{\beta}$, and let $j = \arg\max_{i\in[m]} |\gamma_i|$ and $k = \text{sign}(\gamma_j)$, so that $k\gamma_j \ge \|\xi\|/\sqrt{m}$. Further, choose $\xi \notin B_{2\sqrt{m}\epsilon}$, in which case $\epsilon \le \|\xi\|/(2\sqrt{m})$, and hence

$$\begin{aligned}
&k(\xi - k\epsilon\omega_j^{\beta})'\omega_j^{\beta} = k\gamma_j - \epsilon \ge \|\xi\|/(2\sqrt{m}) \\
\text{and} \quad &\|\xi - k\epsilon\omega_i^{\beta}\| \le 2\|\xi\|.
\end{aligned}$$

Finally, set ϵ small enough so that the $o(\epsilon)$ term in (28) is less than $\max_{i\in[m]} \epsilon/(16\sqrt{m}\sigma_i^{\beta})$, and choose t large enough so that the $O_p(t^{-1/2})$ term is less than $\max_{i\in[m]} \epsilon/(16\sqrt{m}\sigma_i^{\beta})$, with at least three-quarters probability. When this last event happens, the previous two inequalities and line

(28) yield the following, for all $\xi \notin B_{2\sqrt{m}\epsilon}(0)$:

$$\begin{aligned}
\bar{\Lambda}_t^{\xi} - \bar{\Lambda}_t^{k\epsilon\omega_j^\beta} &\leq \frac{-nt\epsilon}{n-t}\|\xi\|/(2\sqrt{m}\sigma_j^\beta) + 2t\|\xi\|(o(\epsilon) + O_p(t^{-1/2})) \\
&\leq -t\epsilon\|\xi\|/(2\sqrt{m}\sigma_j^\beta) + 2t\|\xi\|(\epsilon/(16\sqrt{m}\sigma_i^\beta) + \epsilon/(16\sqrt{m}\sigma_j^\beta)) \\
&\leq -t\epsilon\|\xi\|/(4\sqrt{m}\sigma_j^\beta) \\
&\leq 0.
\end{aligned}$$

This establishes line (27).

Fourth, I will use (23) to show that

$$\Pr\left(n\Lambda_n^\beta(y_n^\beta) - \max_{k\in\{-1,1\}}\max_{j\in[m]}\bar{\Lambda}_t^{k\epsilon\omega_j^\beta} \geq n^{2/3}\right) \geq 3/4. \tag{29}$$

This expression implies that there is probably at least one combination of $k \in \{-1,1\}$ and $j \in [m]$ for which allocating $t(\beta + k\epsilon\omega_j^\beta)$ units of inventory to the first t periods is very costly. And with (27), this will imply that there's a decent chance that allocating $t(\beta+\xi)$ units of inventory to the first t periods will be very costly, for *any* $\xi \notin B_{2\sqrt{m}\epsilon}(0)$.

To begin, a nasty series of triangle inequalities yields the following, for $\zeta \equiv k\epsilon\omega_j^\beta$:

$$\begin{aligned}
n\Lambda_n^\beta(y_n^\beta) - \bar{\Lambda}_t^\zeta \geq\ & n\Lambda_\infty^\beta(y_\infty^\beta) - t\Lambda_\infty^{\beta+\zeta}(y_\infty^{\beta+\zeta}) + (n-t)\Lambda_\infty^{\beta-\frac{t}{n-t}\zeta}(y_\infty^{\beta-\frac{t}{n-t}\zeta}) \\
& - n|\Lambda_\infty^\beta(y_n^\beta) - \Lambda_\infty^\beta(y_\infty^\beta)| \\
& - t|\Lambda_\infty^{\beta+\zeta}(y_t^{\beta+\zeta}) - \Lambda_\infty^{\beta+\zeta}(y_\infty^{\beta+\zeta})| \\
& - (n-t)|\Lambda_\infty^{\beta-\frac{t}{n-t}\zeta}(\underleftarrow{y}_t^{\beta-\frac{t}{n-t}\zeta}) - \Lambda_\infty^{\beta-\frac{t}{n-t}\zeta}(y_\infty^{\beta-\frac{t}{n-t}\zeta})| \\
& - t\left|\Lambda_t^\beta(y_n^\beta) - \Lambda_\infty^\beta(y_n^\beta) - \Lambda_t^{\beta+\zeta}(y_t^{\beta+\zeta}) + \Lambda_\infty^{\beta+\zeta}(y_t^{\beta+\zeta})\right| \\
& - (n-t)\left|\underleftarrow{\Lambda}_t^\beta(y_n^\beta) - \Lambda_\infty^\beta(y_n^\beta) - \underleftarrow{\Lambda}_t^{\beta-\frac{t}{n-t}\zeta}(\underleftarrow{y}_t^{\beta-\frac{t}{n-t}\zeta}) + \Lambda_\infty^{\beta-\frac{t}{n-t}\zeta}(\underleftarrow{y}_t^{\beta-\frac{t}{n-t}\zeta})\right|.
\end{aligned}$$

Each term on the right is bounded in probability: First, line (23) establishes that $n\Lambda_\infty^\beta(y_\infty^\beta) - t\Lambda_\infty^{\beta+\zeta}(y_\infty^{\beta+\zeta}) + (n-t)\Lambda_\infty^{\beta-\frac{t}{n-t}\zeta}(y_\infty^{\beta-\frac{t}{n-t}\zeta}) = \Omega(t)$. Second, since Λ_∞^β is differentiable and since $\|y_n^\beta - y_\infty^\beta\| = O_p(n^{-1/2})$, by Proposition 4, we have $n|\Lambda_\infty^\beta(y_n^\beta) - \Lambda_\infty^\beta(y_\infty^\beta)| = O_p(n^{1/2})$. Likewise, $t|\Lambda_\infty^{\beta+\zeta}(y_t^{\beta+\zeta}) - \Lambda_\infty^{\beta+\zeta}(y_\infty^{\beta+\zeta})|$ and $(n-t)|\Lambda_\infty^{\beta-\frac{t}{n-t}\zeta}(\underleftarrow{y}_t^{\beta-\frac{t}{n-t}\zeta}) - \Lambda_\infty^{\beta-\frac{t}{n-t}\zeta}(y_\infty^{\beta-\frac{t}{n-t}\zeta})|$ are $O_p(t^{1/2})$ and $O_p((n-t)^{1/2})$, respectfully. Finally, Proposition 4 and Lemma 7 imply that the last two terms are also $O_p(t^{1/2})$ and $O_p((n-t)^{1/2})$. Accordingly, we can set n large enough so that if $n^{3/4} \leq t \leq n/2$

then there is at least a 75% chance that (i) the $\Omega(t)$ term exceeds $2n^{2/3}$ and (ii) the sum of the $O_p(n^{1/2})$, $O_p(t^{1/2})$, and $O_p((n-t)^{1/2})$ terms are no more than $n^{2/3}$, for all combinations of $k \in \{-1, 1\}$ and $j \in [m]$. And this establishes (29).

Finally, combining (22), (27), and (29) yields the following, for sufficiently small ϵ, sufficiently large n, and $n^{3/4} \le t \le n/2$:

$$\begin{aligned}
&\Pr(\hat{R}_n \ge n^{2/3} \mid b_t \notin B_{2\sqrt{m}\epsilon}(\beta)) \\
&\quad \ge \Pr(n\Lambda_n^\beta(y_n^\beta) - \bar{\Lambda}_t^{b_t-\beta} \ge n^{2/3} \mid b_t \notin B_{2\sqrt{m}\epsilon}(\beta)) \\
&\quad \ge \Pr\big(n\Lambda_n^\beta(y_n^\beta) - \sup_{\xi \notin B_{2\sqrt{m}\epsilon}(0)} \bar{\Lambda}_t^\xi \ge n^{2/3}\big) \\
&\quad \ge \Pr\big(n\Lambda_n^\beta(y_n^\beta) - \max_{k\in\{-1,1\}} \max_{j\in[m]} \bar{\Lambda}_t^{k\epsilon\omega_j^\beta} \ge n^{2/3} \ \cap \sup_{\xi \notin B_{2\sqrt{m}\epsilon}(0)} \bar{\Lambda}_t^\xi \le \max_{k\in\{-1,1\}} \max_{j\in[m]} \bar{\Lambda}_t^{k\epsilon\omega_j^\beta}\big) \\
&\quad \ge \Pr\big(n\Lambda_n^\beta(y_n^\beta) - \max_{k\in\{-1,1\}} \max_{j\in[m]} \bar{\Lambda}_t^{k\epsilon\omega_j^\beta} \ge n^{2/3}\big) + \Pr\big(\sup_{\xi \notin B_{2\sqrt{m}\epsilon}(0)} \bar{\Lambda}_t^\xi \le \max_{k\in\{-1,1\}} \max_{j\in[m]} \bar{\Lambda}_t^{k\epsilon\omega_j^\beta}\big) - 1 \\
&\quad \ge 3/4 + 3/4 - 1 \\
&\quad = 1/2.
\end{aligned}$$

Finally, since $\mathrm{E}(\hat{R}_n) = O(\log n)$, by Theorem 1 and line (33), the result above implies

$$\begin{aligned}
O(\log n) &= \mathrm{E}(\hat{R}_n) \\
&\ge n^{2/3} \Pr(b_t \notin B_{2\sqrt{m}\epsilon}(\beta)) \Pr(\hat{R}_n \ge n^{2/3} \mid b_t \notin B_{2\sqrt{m}\epsilon}(\beta)) \\
&\ge n^{2/3} \Pr(b_t \notin B_{2\sqrt{m}\epsilon}(\beta))/2.
\end{aligned}$$

And this implies the result. □

Lemma 10 Proof. First, note that $\pi_t^{b_t} = 1$ implies $b_{t-1} = \psi_t^{b_t}(a_t)$, and hence $v_{t-1}^{\psi_t^{b_t}(a_t)} = \bar{V}_{t-1}^{\psi_t^{b_t}(a_t)} -$

$\hat{R}_{t-1}$. Also, $\pi_t^{b_t} = 1$ implies $tb_t \geq a_t$, which with lines (5) and (6) yield

$$\begin{aligned}
\hat{R}_t &\equiv \bar{V}_t^{b_t} - v_t^{b_t} \\
&= \max_{x_t \in [0,1]} x_t u_t + \bar{V}_{t-1}^{\psi_t^{b_t}(x_t a_t)} - u_t - v_{t-1}^{\psi_t^{b_t}(a_t)} \\
&= \max_{x_t \in [0,1]} (x_t - 1) u_t + \bar{V}_{t-1}^{\psi_t^{b_t}(x_t a_t)} - \bar{V}_{t-1}^{\psi_t^{b_t}(a_t)} + \hat{R}_{t-1} \\
&\geq \max_{x_t \in [0,1]} (x_t - 1) u_t + (1 - x) a_t' y_{t-1}^{\psi_t^{b_t}(0)} + \hat{R}_{t-1} \\
&= (a_t' y_{t-1}^{\psi_t^{b_t}(0)} - u_t)^+ + \hat{R}_{t-1} \\
&= \Delta_t (y_{t-1}^{\psi_t^{b_t}(0)})^- + \hat{R}_{t-1}.
\end{aligned}$$

Analogously, if $\pi_t^{b_t} = 0$ and $tb_t \geq a_t$ then lines (5) and (7) yield

$$\hat{R}_t \geq \Delta_t (y_{t-1}^{\psi_t^{b_t}(a_t)})^+ + \hat{R}_{t-1}.$$

Further, we always trivially have

$$\hat{R}_t \geq \hat{R}_{t-1}.$$

Now choose $\delta > 0$ small enough to ensure that $\delta\iota \leq \beta$, where ι is a vector of ones. In this case, $b_t \in B_{\delta/2}(\beta)$ implies $tb_t \geq a_t$ for $t \geq 2\|\alpha\|/\delta$, which with our previous three inequalities inductively yields the result. □

Lemma 11 Proof. I will show that there exists $C > 0$ that satisfies

$$\inf_{b \in B_{\delta/2}(\beta)} \mathrm{E}\left(\pi_t^b \Delta_t (y_{t-1}^{\psi_t^b(0)})^- + (1 - \pi_t^b) \Delta_t (y_{t-1}^{\psi_t^b(\alpha)})^+\right) \geq C m \sigma_m^\beta / (2t), \tag{30}$$

for all sufficiently large t. Combining this result with Lemma 9 yields the desired result:

$$\begin{aligned}
\mathrm{E}(r_t) &= \mathrm{E}\left(\mathbb{1}\{b_t \in B_{\delta/2}(\beta)\}\left(\pi_t^{b_t} \Delta_t (y_{t-1}^{\psi_t^{b_t}(0)})^- + (1 - \pi_t^{b_t}) \Delta_t (y_{t-1}^{\psi_t^{b_t}(a_t)})^+\right)\right) \\
&= \mathrm{E}\left(\mathbb{1}\{b_t \in B_{\delta/2}(\beta)\}\, \mathrm{E}\left(\pi_t^b \Delta_t (y_{t-1}^{\psi_t^b(0)})^- + (1 - \pi_t^b) \Delta_t (y_{t-1}^{\psi_t^b(a_t)})^+\right)\Big|_{b = b_t}\right) \\
&\geq \Pr(b_t \in B_{\delta/2}(\beta)) \inf_{b \in B_{\delta/2}(\beta)} \mathrm{E}\left(\pi_t^b \Delta_t (y_{t-1}^{\psi_t^b(0)})^- + (1 - \pi_t^b) \Delta_t (y_{t-1}^{\psi_t^b(\alpha)})^+\right) \\
&\geq (1 - n^{-1/2}) C m \sigma_m^\beta / (2t),
\end{aligned}$$

where the second line holds because b_t is independent of the random mapping $b \mapsto \pi_t^b \Delta_t(y_{t-1}^{\psi_t^b(0)})^- + (1-\pi_t^b)\Delta_t(y_{t-1}^{\psi_t^b(a_t)})^+$, and the last line holds because t satisfies $n^{3/4} \le t \le n/2$.

Let me briefly outline how we will establish line (30). First, Lemma 2 implies that there's a $O(1)$ chance that we underestimate the shadow price by at least $4\iota/\sqrt{t}$. I use this fact to establish that there exists some constant $C > 0$ that satisfies the following for sufficiently large t:

$$
\begin{aligned}
&\mathrm{E}\left((1-\pi_t^b)\Delta_t(y_{t-1}^{\psi_t^b(a_t)})^+\right) \\
&\quad \ge C\,\mathrm{E}\left((1-\pi_t^b)\big(\mathbb{1}\{\Delta_t(y_\infty^b - \iota/\sqrt{t}) > 0\} - \mathbb{1}\{\Delta_t(y_\infty^b + \iota/\sqrt{t}) > 0\}\big)a_t'\iota/\sqrt{t}\right).
\end{aligned}
$$

This lower bound looks nasty, but it's almost exactly in the form we need to apply our one remaining tool: Assumption 6. However, before applying this assumption, I must eliminate the pesky $1-\pi_t^b$ term. Fortunately, $\mathrm{E}\left(\pi_t^b\Delta_t(y_{t-1}^{\psi_t^b(0)})^-\right)$ honors the same bound, except with π_t^b replacing $1-\pi_t^b$, which means that $\mathrm{E}\left(\pi_t^b\Delta_t(y_{t-1}^{\psi_t^b(0)})^- + (1-\pi_t^b)\Delta_t(y_{t-1}^{\psi_t^b(\alpha)})^+\right)$ has a corresponding π_t^b-free bound, which makes it amenable to Assumption 6. Finally, the last part of the proof combines this assumption with the fundamental theorem of calculus and Lemma 1 to express this expectation as an integral over $\ddot{\Lambda}_\infty$.

To begin the proof, let t be large enough so that $b \in B_{\delta/2}(\beta)$ implies $\psi_t^b(a_t) \in B_\delta(\beta)$. In this case Lemma 2 establishes that there exists $C > 0$ that satisfies the following, for all $b \in B_{\delta/2}(\beta)$:

$$
\begin{aligned}
&\Pr\left(\|\sqrt{t}(y_{t-1}^{\psi_t^b(a_t)} - y_\infty^{\psi_t^b(a_t)}) + 4\iota\| \le 1\right) \\
&\quad \ge \Pr\Big(\sup_{b\in B_\delta(\beta)} \|\sqrt{t}(y_{t-1}^b - y_\infty^b) + 4\iota\| \le 1\Big) \\
&\quad = \Pr\Big(\sup_{b\in B_\delta(\beta)} \|\sqrt{t}(y_{t-1}^b - y_\infty^b) + 4\iota\| \le 1\Big) \\
&\quad > C.
\end{aligned}
$$

Furthermore, $\|\sqrt{t}(y_{t-1}^{\psi_t^b(a_t)} - y_\infty^{\psi_t^b(a_t)}) + 4\iota\| \le 1$ implies the following, when t is large:

$$
\begin{aligned}
y_{t-1}^{\psi_t^b(a_t)} - y_\infty^b &= (y_{t-1}^{\psi_t^b(a_t)} - y_\infty^{\psi_t^b(a_t)}) + (y_\infty^{\psi_t^b(a_t)} - y_\infty^b) \\
&\le -3\iota/\sqrt{t} + \iota/\sqrt{t} \\
&= -2\iota/\sqrt{t},
\end{aligned}
$$

where the second line follows because $\|y_\infty^{\psi_t^b(a_t)} - y_\infty^b\| = o(1/\sqrt{t})$, by Assumption 4 and Lemma 2.

Now combining the previous two results yields the following, for $b \in B_{\delta/2}(\beta)$ and t large:

$$\begin{aligned}
&\mathrm{E}\left((1-\pi_t^b)\Delta_t(y_{t-1}^{\psi_t^b(a_t)})^+\right) \\
&\quad\geq \mathrm{E}\left(\mathbb{1}\{\|\sqrt{t}(y_{t-1}^{\psi_t^b(a_t)} - y_\infty^{\psi_t^b(a_t)}) + 4\iota\| \leq 1\}(1-\pi_t^b)\Delta_t(y_\infty^b - 2\iota/\sqrt{t})^+\right) \\
&\quad\geq C\,\mathrm{E}\left((1-\pi_t^b)\Delta_t(y_\infty^b - 2\iota/\sqrt{t})^+\right) \\
&\quad\geq C\,\mathrm{E}\left((1-\pi_t^b)\mathbb{1}\{\Delta_t(y_\infty^b - \iota/\sqrt{t}) \geq 0\}\Delta_t(y_\infty^b - 2\iota/\sqrt{t})^+\right) \\
&\quad\geq C\,\mathrm{E}\left((1-\pi_t^b)\mathbb{1}\{\Delta_t(y_\infty^b - \iota/\sqrt{t}) \geq 0\}a_t'\iota/\sqrt{t}\right) \\
&\quad\geq C\,\mathrm{E}\left((1-\pi_t^b)\big(\mathbb{1}\{\Delta_t(y_\infty^b - \iota/\sqrt{t}) > 0\} - \mathbb{1}\{\Delta_t(y_\infty^b + \iota/\sqrt{t}) > 0\}\big)a_t'\iota/\sqrt{t}\right).
\end{aligned}$$

Note, the third line above holds because $y_{t-1}^{\psi_t^b(a_t)}$ is independent of Δ_t and π_t, and the fifth line holds because $\Delta_t(y_\infty^b - \iota/\sqrt{t}) \geq 0$ implies $u_t \geq a_t' y_\infty^b - a_t'\iota/\sqrt{t}$ and hence implies $\Delta_t(y_\infty^b - 2\iota/\sqrt{t})^+ \geq a_t'\iota/\sqrt{t}$.

Next, an analogous argument implies that we can set C small enough to satisfy the following, for $b \in B_{\delta/2}(\beta)$ and large t:

$$\begin{aligned}
&\mathrm{E}\left(\pi_t^b\Delta_t(y_{t-1}^{\psi_t^b(0)})^-\right) \\
&\quad\geq C\,\mathrm{E}\left(\pi_t^b\big(\mathbb{1}\{\Delta_t(y_\infty^b - \iota/\sqrt{t}) > 0\} - \mathbb{1}\{\Delta_t(y_\infty^b + \iota/\sqrt{t}) > 0\}\big)a_t'\iota/\sqrt{t}\right).
\end{aligned}$$

Finally, adding our two bounds establishes line (30):

$$\begin{aligned}
&\mathrm{E}\left(\pi_t^b\Delta_t(y_{t-1}^{\psi_t^b(0)})^- + (1-\pi_t^b)\Delta_t(y_{t-1}^{\psi_t^b(\alpha)})^+ \mid b \in B_{\delta/2}(\beta)\right) \\
&\quad= C\,\mathrm{E}\left(\big(\mathbb{1}\{\Delta_t(y_\infty^b - \iota/(2\sqrt{t})) > 0\} - \mathbb{1}\{\Delta_t(y_\infty^b + \iota/(2\sqrt{t})) > 0\}\big)a_t'\iota/\sqrt{t}\right) \\
&\quad= C/\sqrt{t}\int_{\gamma=-1}^{1} \iota' \tfrac{\partial}{\partial\gamma}\,\mathrm{E}\left(\mathbb{1}\{\Delta_1(y_\infty^b - \gamma\iota/(2\sqrt{t})) > 0\}a_1\right)d\gamma \\
&\quad= C/\sqrt{t}\int_{\gamma=-1}^{1} \iota'\ddot{\Lambda}_\infty(y_\infty^b - \gamma\iota/\sqrt{t})\iota/(2\sqrt{t})d\gamma \\
&\quad\geq C/\sqrt{t}\int_{\gamma=-1}^{1} \iota'\iota\sigma_m^\beta/(4\sqrt{t})d\gamma \\
&\quad\geq Cm\sigma_m^\beta/(2t).
\end{aligned}$$

The penultimate line above holds because the smallest singular value of $\ddot{\Lambda}_\infty(y_\infty^b - \gamma\iota/\sqrt{t})$ is at least half of the smallest singular value of $\ddot{\Lambda}_\infty(y_\infty^\beta)$, when b is near β and t is large. □

Lemma 1. $(y_t^b - y)'\dot{\Lambda}_t^b(y) \leq 0$ *for all* $t \in \mathbb{N}$, $b \in \mathbb{R}_+^m$, *and* $y \in \mathbb{R}_+^m$.

Proof. Since $\dot{\Lambda}_t^b$ is a subgradient, it satisfies $\Lambda_t^b(y_t^b) - \Lambda_t^b(y) \geq (y_t^b - y)'\dot{\Lambda}_t^b(y)$. And since $\Lambda_t^b(y_t^b) \leq$

$\Lambda_t^b(y)$, this yields the result. $\square$

Lemma 2. *For all* $\gamma \in \mathbb{R}^m$ *and* $\epsilon > 0$ *there exist* $\delta, C > 0$ *such that* $\Pr(\sup_{b \in B_\delta(\beta)} \|\sqrt{t}(y_t^b - y_\infty^b) - \gamma\| \leq \epsilon) \geq C$ *for all sufficiently large* t.

Proof. I will begin with a brief proof sketch. Our primary tool for positioning y_t^b is is Lemma 3, which maintains that y_t^b will be close to $y_\infty^b + \gamma/\sqrt{t}$ for all $b \in B_\delta(\beta)$ if $\dot{\Lambda}_t^b(y_\infty^b + (\gamma + \eta k \omega_j^b)/\sqrt{t})$ is close to $\eta k \sigma_j^b \omega_j^b$, for all $b \in B_\delta(\beta)$, $j \in [m]$, and $k \in \{-1, 1\}$. And with a few triangle inequalities and some basic calculus, I show that this condition holds when $\sqrt{t}(\dot{\Lambda}_t^\beta(y) - \dot{\Lambda}_\infty^\beta(y))$ is near $\ddot{\Lambda}_\infty(y_\infty^\beta)\gamma$ for all y in the ν-ball of y_∞^β, for some $\nu > 0$. Finally, I use Lemma 8 to show that there's an $\Theta(1)$ chance of this happening. This lemma maintains that the mapping $(j, y) \mapsto \sqrt{t} e_j'(\dot{\Lambda}_t^\beta(y) - \dot{\Lambda}_\infty^\beta(y))$ converges to a Gaussian process whose mean is near $\ddot{\Lambda}_\infty(y_\infty^\beta)\gamma$ when we condition on $\sqrt{t}(\dot{\Lambda}_t^\beta(y_\infty^\beta) - \dot{\Lambda}_\infty^\beta(y_\infty^\beta))$ being near $\ddot{\Lambda}_\infty(y_\infty^\beta)\gamma$.

Now, I will begin the proof in earnest. Lemmas 1 and 2 imply that we can choose δ small enough so that $\sigma_1^b \leq 2\sigma_1^\beta$ and $\sigma_m^b \geq \sigma_m^\beta/2 > 0$ for all $b \in B_\delta(\beta)$. And these lemmas also imply that we can choose t large enough to ensure that $y_\infty^b + (\gamma + \eta k \omega_j^b)/\sqrt{t} \geq 0$ for a given $\eta > 0$ and all $j \in [m]$, $k \in \{-1, 1\}$, and $b \in B_\delta(\beta)$. With this, Lemma 3 indicates that $\sup_{b \in B_\delta(\beta)} \|\sqrt{t}(y_t^b - y_\infty^b) - \gamma\| \leq \epsilon$ if

$$\sup_{b \in B_\delta(\beta)} \max_{j \in [m]} \max_{k \in \{-1,1\}} \|\sqrt{t}\dot{\Lambda}_t^b(y_\infty^b + (\gamma + \eta k \omega_j^b)/\sqrt{t}) - \eta k \sigma_j^b \omega_j^b\| \leq \kappa, \tag{31}$$

where $\eta \equiv \epsilon/(1 + 8\sqrt{m}\sigma_1^\beta/\sigma_m^\beta)$ and $\kappa \equiv \eta\sigma_m^\beta/(4\sqrt{m})$. Further, this inequality holds when the following inequalities hold for all $b \in B_\delta(\beta)$, $j \in [m]$, and $k \in \{-1, 1\}$:

$$\|\sqrt{t}(\dot{\Lambda}_t^\beta(y_\infty^b + (\gamma + \eta k \omega_j^b)/\sqrt{t}) - \dot{\Lambda}_\infty^\beta(y_\infty^b + (\gamma + \eta k \omega_j^b)/\sqrt{t})) + \ddot{\Lambda}_\infty(y_\infty^\beta)\gamma\| \leq \kappa/3, \tag{32}$$

$$\|\sqrt{t}\dot{\Lambda}_t^b(y_\infty^b + (\gamma + \eta k \omega_j^b)/\sqrt{t}) - \ddot{\Lambda}_\infty(y_\infty^b)\gamma - \eta k \sigma_j^b \omega_j^b\| \leq \kappa/3, \tag{33}$$

$$\text{and} \quad \|\ddot{\Lambda}_\infty(y_\infty^b)\gamma - \ddot{\Lambda}_\infty(y_\infty^\beta)\gamma\| \leq \kappa/3. \tag{34}$$

Lines (32)–(34) imply line (31) by the triangle inequality, and by the fact that $\dot{\Lambda}_t^b(y) - \dot{\Lambda}_\infty^b(y)$ is independent of b, which enables me to change the superscripts in line (32) from b to β. I will now show that there's a non-negligible chance that these inequalities hold universally across b, j, and k in their respective domains.

First, since $B_\delta(\beta)$ is compact and $\ddot{\Lambda}_\infty(y_\infty^b)$ is continuous in b, by Lemmas 1 and 2, it follows that we can set δ small enough to make inequality (34) hold universally.

Second, since $\dot{\Lambda}^b_\infty(y^b_\infty) = 0$ and $\ddot{\Lambda}_\infty$ is locally continuous near y^β_∞, the mean value theorem indicates that there exists $\xi^b_{tjk} \in (0,1)$ for which

$$\begin{aligned}
&\sqrt{t}\dot{\Lambda}^b_\infty(y^b_\infty + (\gamma + \eta k \omega^b_j)/\sqrt{t}) \\
&\qquad = \sqrt{t}(\dot{\Lambda}^b_\infty(y^b_\infty + (\gamma + \eta k \omega^b_j)/\sqrt{t}) - \dot{\Lambda}^b_\infty(y^b_\infty)) \\
&\qquad = \ddot{\Lambda}_\infty(y^b_\infty + \xi^b_{tjk}(\gamma + \eta k \omega^b_j)/\sqrt{t})(\gamma + \eta k \omega^b_j) \\
&\qquad = \ddot{\Lambda}_\infty(y^b_\infty)\gamma + \eta k \sigma^b_j \omega^b_j + \zeta^b_{tjk}(\gamma + \eta k \omega^b_j), \\
\text{where} \quad &\zeta^b_{tjk} \equiv \ddot{\Lambda}_\infty(y^b_\infty + \xi^b_{tjk}(\gamma + \eta k \omega^b_j)/\sqrt{t}) - \ddot{\Lambda}_\infty(y^b_\infty).
\end{aligned}$$

And the continuity of $\ddot{\Lambda}_\infty$ implies that we can set δ small enough so that

$$\sup_{b \in B_\delta(\beta)} \max_{j \in [m]} \max_{k \in \{-1,1\}} \|\zeta^b_{tjk}(\gamma + \eta k \omega^b_j)\| \leq \kappa/3,$$

for all sufficiently large t. Hence, inequality (33) will hold universally for all sufficiently large t and small δ.

Finally, I will show that for all sufficiently large t the probability that inequality (32) holds universally across $b \in B_\delta(\beta)$, $j \in [m]$, and $k \in \{-1,1\}$ exceeds some $C > 0$. Since y^b_∞ and ω^b_j are continuous in b, by Lemmas 1 and 2, it will suffice to show that there exist $\nu > 0$ such that

$$\liminf_{t \to \infty} \Pr\Big(\sup_{y \in B_\nu(y^\beta_\infty)} \|\sqrt{t}(\dot{\Lambda}^\beta_t(y) - \dot{\Lambda}^\beta_\infty(y)) + \ddot{\Lambda}_\infty(y^\beta_\infty)\gamma\| \leq \kappa/3\Big) > 0.$$

I will prove this inequality with Lemma 8, which with proposition 3.13 of Eaton (1983) implies that conditional on $\zeta_t \equiv \sqrt{t}(\dot{\Lambda}^\beta_t(y^\beta_\infty) - \dot{\Lambda}^\beta_\infty(y^\beta_\infty))$, the random map $(j, y) \mapsto \sqrt{t}e'_j(\dot{\Lambda}^\beta_t(y) - \dot{\Lambda}^\beta_\infty(y))$ weakly converges to a Gaussian process with domain $[m] \times B_\nu(y^\beta_\infty)$, mean function ρ^{ζ_t}, and covariance function Ξ, where

$$\begin{aligned}
\rho^{\zeta_t}_j(y) &\equiv e'_j\Omega(y, y^\beta_\infty)\Omega(y^\beta_\infty, y^\beta_\infty)^{-1}\zeta_t, \\
\Xi_{j\bar{j}}(y, \bar{y}) &\equiv e'_j\Omega(y, \bar{y})e_{\bar{j}} - e'_j\Omega(y, y^\beta_\infty)\Omega(y^\beta_\infty, y^\beta_\infty)^{-1}\Omega(y^\beta_\infty, \bar{y})e_{\bar{j}}, \\
\text{and} \quad \Omega(y, \bar{y}) &\equiv \mathrm{E}(\mathbb{1}\{\Delta_1(y) > 0\}\mathbb{1}\{\Delta_1(\bar{y}) > 0\}a_1a'_1) - \mathrm{E}(\mathbb{1}\{\Delta_1(y) > 0\}a_1)\,\mathrm{E}(\mathbb{1}\{\Delta_1(\bar{y}) > 0\}a'_1).
\end{aligned}$$

Since Ξ is independent of ζ_t, the random map $(j, y) \mapsto \sqrt{t}e'_j(\dot{\Lambda}^\beta_t(y) - \dot{\Lambda}^\beta_\infty(y)) - e'_j\rho^{\zeta_t}_j(y)$ is assymptotically independent of ζ_t, and hence assymptotically independent of the random map $y \mapsto \rho^{\zeta_t}(y)$.

Accordingly, for sufficiently large t, we have

$$
\begin{aligned}
&\Pr\Big(\sup_{y\in B_\nu(y^\beta_\infty)} \|\sqrt{t}(\dot\Lambda^\beta_t(y)-\dot\Lambda^\beta_\infty(y))+\ddot\Lambda_\infty(y^\beta_\infty)\gamma\| \le \kappa/3\Big)\\
&\quad\ge \Pr\Big(\sup_{y\in B_\nu(y^\beta_\infty)} \|\sqrt{t}(\dot\Lambda^\beta_t(y)-\dot\Lambda^\beta_\infty(y))-\rho^{\zeta_t}(y)\| \le \kappa/6\\
&\qquad\qquad \cap \sup_{y\in B_\nu(y^\beta_\infty)} \|\rho^{\zeta_t}(y)-\ddot\Lambda_\infty(y^\beta_\infty)\gamma\| \le \kappa/6\Big)\\
&\quad\ge p^1_t p^2_t/2,\\
\text{where}\quad & p^1_t \equiv \Pr\big(\sup_{y\in B_\nu(y^\beta_\infty)} \|\sqrt{t}(\dot\Lambda^\beta_t(y)-\dot\Lambda^\beta_\infty(y))-\rho^{\zeta_t}(y)\| \le \kappa/6\big)\\
\text{and}\quad & p^2_t \equiv \Pr\big(\sup_{y\in B_\nu(y^\beta_\infty)} \|\rho^{\zeta_t}(y)-\ddot\Lambda_\infty(y^\beta_\infty)\gamma\| \le \kappa/6\big).
\end{aligned}
$$

I will now lower bound probability p^2_t. It is straightforward to confirm that $\|\zeta_t-\rho^{\zeta_t}(y)\| = O(\|y-y^\beta_\infty\|)O(\|\zeta_t\|)$, which implies that we can choose ν small enough so that $\|\zeta_t-\ddot\Lambda_\infty(y^\beta_\infty)\gamma\| \le \kappa/12$ implies $\|\zeta_t-\rho^{\zeta_t}(y)\| \le \kappa/12$ for all $y\in B_\nu(y^\beta_\infty)$. This, in turn, implies that

$$
\begin{aligned}
p^2_t &\ge \Pr\Big(\|\zeta_t-\ddot\Lambda_\infty(y^\beta_\infty)\gamma\| \le \kappa/12 \ \cap \sup_{y\in B_\nu(y^\beta_\infty)} \|\zeta_t-\rho^{\zeta_t}(y)\| \le \kappa/12\Big)\\
&= \Pr\Big(\|\zeta_t-\ddot\Lambda_\infty(y^\beta_\infty)\gamma\| \le \kappa/12\Big).
\end{aligned}
$$

Finally, the limit inferior of this last probability is strictly positive, as $t\to\infty$, because ζ_t converges to a multivariate normal with a full-rank covariance matrix, by Lemma 8.

I will now lower bound probability p^1_t. First, $\|\zeta_t-\rho^{\zeta_t}(y)\| = O(\|y-y^\beta_\infty\|)O(\|\zeta_t\|)$ implies that for a given $M>0$ we can set ν small enough so that $\|\zeta_t\|\le M$ implies $\|\zeta_t-\rho^{\zeta_t}(y)\| \le \kappa/12$ for all $y\in B_\nu(y^\beta_\infty)$. And since ζ_t converges to a multivariate normal, we can choose M large enough so that the last equality below holds for all sufficiently large t:

$$
\begin{aligned}
&\Pr\big(\sup_{y\in B_\nu(y^\beta_\infty)} \|\sqrt{t}(\dot\Lambda^\beta_t(y)-\dot\Lambda^\beta_\infty(y))-\rho^{\zeta_t}(y)\| \le \kappa/6\big)\\
&\quad\ge \Pr\big(\sup_{y\in B_\nu(y^\beta_\infty)} \|\sqrt{t}(\dot\Lambda^\beta_t(y)-\dot\Lambda^\beta_\infty(y))-\zeta_t\| \le \kappa/12 \ \cap \sup_{y\in B_\nu(y^\beta_\infty)} \|\zeta_t-\rho^{\zeta_t}(y)\| \le \kappa/12\big)\\
&\quad\ge \Pr\big(\sup_{y\in B_\nu(y^\beta_\infty)} \|\sqrt{t}(\dot\Lambda^\beta_t(y)-\dot\Lambda^\beta_\infty(y))-\zeta_t\| \le \kappa/12 \ \cap \ \|\zeta_t\| < M\big)\\
&\quad\ge \Pr\big(\sup_{y\in B_\nu(y^\beta_\infty)} \|\sqrt{t}(\dot\Lambda^\beta_t(y)-\dot\Lambda^\beta_\infty(y))-\zeta_t\| \le \kappa/12\big)/2.
\end{aligned}
$$

Further, Lemma 6 implies that we can set ν small enough so that the last inequality in the expression below holds:

$$\begin{aligned}
&\Pr\big(\sup_{y\in B_\nu(y^\beta_\infty)} \|\sqrt{t}(\dot\Lambda^\beta_t(y)-\dot\Lambda^\beta_\infty(y))-\zeta_t\| > \kappa/12\big)\\
&\quad= \Pr\big(\sup_{y\in B_\nu(y^\beta_\infty)} \|\sqrt{t}(\dot\Lambda^\beta_t(y)-\dot\Lambda^\beta_\infty(y))-\zeta_t\|^2 > \kappa^2/144\big)\\
&\quad\le \mathrm{E}\big(\sup_{y\in B_\nu(y^\beta_\infty)} \|\sqrt{t}(\dot\Lambda^\beta_t(y)-\dot\Lambda^\beta_\infty(y))-\zeta_t\|^2\big)/(\kappa^2/144)\\
&\quad\le 1/2.
\end{aligned}$$

Accordingly, we can set ν small enough so that $p^1_t \ge 1/4$ for all sufficiently large t. ☐

Lemma 3. *If b is close enough to β to ensure that $\{\omega^b_i\}_{i\in[m]}$ and $\{\sigma^b_i\}_{i\in[m]}$ exist, and if $y\in\mathbb{R}^m_{>0}$ and $\eta>0$ satisfy $y+\eta k\omega^b_j\ge 0$ and $\|\dot\Lambda^b_t(y+\eta k\omega^b_j)-\eta k\sigma^b_j\omega^b_j\|\le \eta\sigma^b_m/(2\sqrt{m})$ for all $j\in[m]$ and $k\in\{-1,1\}$ then $y^b_t\in B_{\eta(1+2\sqrt{m}\sigma^b_1/\sigma^b_m)}(y)$.*

Proof. Combining Lemma 1 with the hypotheses of the current lemma implies the following:

$$\begin{aligned}
0 &\ge (y^b_t-y-\eta k\omega^b_j)'\dot\Lambda^b_t(y+\eta k\omega^b_j)\\
&= (y^b_t-y-\eta k\omega^b_j)'\eta k\sigma^b_j\omega^b_j + (y^b_t-y-\eta k\omega^b_j)'(\dot\Lambda^b_t(y+\eta k\omega^b_j)-\eta k\sigma^b_j\omega^b_j)\\
&\ge \eta k\sigma^b_j(y^b_t-y)'\omega^b_j - \eta^2k^2\sigma^b_j\omega^{b\prime}_j\omega^b_j - \|y^b_t-y-\eta k\omega^b_j\|\|\dot\Lambda^b_t(y+\eta k\omega^b_j)-\eta k\sigma^b_j\omega^b_j\|\\
&\ge \eta k\sigma^b_m(y^b_t-y)'\omega^b_j - \eta^2\sigma^b_1 - (\|y^b_t-y\|+\eta)\eta\sigma^b_m/(2\sqrt{m}).
\end{aligned}$$

And since $\omega^b_1,\cdots,\omega^b_m$ are orthonormal, there must be at least one $j\in[m]$ and one $k\in\{-1,1\}$ for which $k(y^b_t-y)'\omega^b_j\ge\|y^b_t-y\|/\sqrt{m}$. And thus, we must have

$$0\ge \eta\sigma^b_m\|y^b_t-y\|/\sqrt{m}-\eta^2\sigma^b_1-\big(\|y^b_t-y\|+\eta\big)\eta\sigma^b_m/(2\sqrt{m}).$$

Finally, rearranging the terms yields the result. ☐

Lemma 4. *There exists $\epsilon>0$ such that if $y,\bar y\in B_\epsilon(y^\beta_\infty)$ then $\mathrm{E}(\mathbb{1}\{\Delta_1(y)>0\}\Delta_1(\bar y)^-)\le 2\sigma^\beta_1\|\bar y-y\|^2$.*

Proof. I will first consider the case in which $\bar y\ge y$. To begin, note that $\mathbb{1}\{\Delta_1(y+dy)>0\}\ne\mathbb{1}\{\Delta_1(y)>0\}$ implies that $u_1=a_1'(y+O(dy))$, in which case $\Delta_1(\bar y)^-=a_1'(\bar y-y+O(dy))$. And with this, Assumption 6 and Lemma 1 imply the following, for y near y_∞^β:

$$\begin{aligned}
&\mathrm{E}\left(\left(\mathbb{1}\{\Delta_1(y+dy)>0\}-\mathbb{1}\{\Delta_1(y)>0\}\right)\Delta_1(\bar{y})^-\right)\\
&\quad=\mathrm{E}\left(\left(\mathbb{1}\{\Delta_1(y+dy)>0\}-\mathbb{1}\{\Delta_1(y)>0\}\right)a_1\right)'(\bar{y}-y+O(dy))\\
&\quad=\left(\tfrac{\partial}{\partial y}\mathrm{E}(\mathbb{1}\{\Delta_1(y)>0\}a_1)dy+o(dy)\right)'(\bar{y}-y+O(dy))\\
&\quad=(\bar{y}-y)'\tfrac{\partial}{\partial y}\mathrm{E}(\mathbb{1}\{\Delta_1(y)>0\}a_1)dy+o(dy)\\
&\quad=(y-\bar{y})'\ddot{\Lambda}(y)dy+o(dy).
\end{aligned}$$

Accordingly, for y near y_∞^β we have $\frac{\partial}{\partial y}\mathrm{E}(\mathbb{1}\{\Delta_1(y)>0\}\Delta_1(\bar{y})^-)=(y-\bar{y})'\ddot{\Lambda}(y)$. And thus, for y and $\bar{y}$ sufficiently close to y_∞^β we have

$$\begin{aligned}
&\mathrm{E}(\mathbb{1}\{\Delta_1(y)>0\}\Delta_1(\bar{y})^-)\\
&\quad=\mathrm{E}(\mathbb{1}\{\Delta_1(y)>0\}\Delta_1(\bar{y})^-)-\mathrm{E}(\mathbb{1}\{\Delta_1(\bar{y})>0\}\Delta_1(\bar{y})^-)\\
&\quad=\int_{\gamma=0}^1\tfrac{\partial}{\partial\gamma}\mathrm{E}(\mathbb{1}\{\Delta_1(\bar{y}+\gamma(y-\bar{y}))>0\}\Delta_1(\bar{y})^-\big)d\gamma\\
&\quad=\int_{\gamma=0}^1(y-\bar{y})'\ddot{\Lambda}(\bar{y}+\gamma(y-\bar{y}))(y-\bar{y})d\gamma\\
&\quad\le\int_{\gamma=0}^1 2\sigma_1^\beta\|y-\bar{y}\|^2d\gamma\\
&\quad=2\sigma_1^\beta\|y-\bar{y}\|^2,
\end{aligned}$$

where the penultimate line holds because the largest singular value of $\ddot{\Lambda}(\bar{y}+\gamma(y-\bar{y}))$ is smaller than twice the largest singular value of y_∞^β, when y and $\bar{y}$ are sufficiently close to y_∞^β, by Lemma 1.

Now we can use what we've just established to prove the $\bar{y}\not\geq y$ case, since

$$\begin{aligned}
&\mathrm{E}(\mathbb{1}\{\Delta_1(y)>0\}\Delta_1(\bar{y})^-)\\
&\quad\le\mathrm{E}(\mathbb{1}\{\Delta_1(y)>0\}\Delta_1(\bar{y}\vee y)^-)\\
&\quad\le 2\sigma_1^\beta\|y-\bar{y}\vee y\|^2\\
&\quad\le 2\sigma_1^\beta\|y-\bar{y}\|^2.
\end{aligned}$$

□

Lemma 5. *There exists $C>0$ such that $\mathrm{E}\left(\sup_{b\in\mathbb{R}_+^m}\sup_{y\in\mathbb{R}_+^m}\|\dot{\Lambda}_t^b(y)-\dot{\Lambda}_\infty^b(y)\|^2\right)\le C/t$ for all $t\in\mathbb{N}$.*

Proof. I will show that the conditions of van der Vaart and Wellner's (1996) theorems 2.14.2 and 2.14.5 are satisfied. First, to translate the problem into van der Vaart and Wellner's format, note that

$$
\begin{aligned}
e_j'(\dot{\Lambda}_t^b(y) - \dot{\Lambda}_\infty^b(y)) &= \sum_{s=1}^{t} \lambda_j^y(u_s, a_s)/t - \mathrm{E}(\lambda_j^y(u_1, a_1)), \\
\text{where} \quad \lambda_j^y(u_1, a_1) &\equiv \mathbb{1}\{\Delta_1(y) > 0\} e_j' a_1.
\end{aligned}
$$

The λ_j^y functions lie under an upper envelope—since $|\lambda_j^y(u_1, a_1)| \leq \|\alpha\|$—and so it will suffice to show that van der Vaart and Wellner's (1996) bracketing integral is finite for the set $\{\lambda_j^y\}_{y \in \mathbb{R}_+^m, j \in [m]}$.

To streamline the argument, I will suppose that $a_1 > 0$, almost surely, and that the conditional distribution of u_1 given a_1 is characterized by density function f, which is bounded by some $M \in \mathbb{N}$. These assumptions are not necessary, but the argument is messy without them.

For a given $\nu > 0$, define m-dimensional grid $G \equiv \gamma \mathbb{Z}^m$, where $\gamma \equiv \nu / (M \|\alpha\|_1^2)$. Next, let $\ell(y) \equiv \max\{g \in G \mid g \leq y\}$ represent the largest gridpoint that's weakly less than $y \in \mathbb{R}_+^m$ and let $h(y) \equiv \min\{g \in G \mid g > \ell(y)\}$ represent the smallest gridpoint that's strictly larger than $\ell(y)$, so that $h(y) - \ell(y) = \gamma \iota$. Finally, define $U \equiv F_u^{-1}(1 - \nu/(2\|\alpha\|))$, $A \equiv F_a^{-1}(\nu/(2\|\alpha\|))$, and $Y \equiv U/A$, where F_u is the CDF of u_1 and F_a is the CDF of the smallest element of a_1 (which we've assumed to be larger than zero).

I will now show that the pair $(\mathbb{1}\{\|y\|_\infty \leq Y\} \lambda_j^{h(y)}, \lambda_j^{\ell(y \wedge Y)})$ is a ν-bracket that contains λ_j^y. First, if $\|y\|_\infty \leq Y$ then

$$
\begin{aligned}
&\mathrm{E}\left((\lambda_j^{\ell(y \wedge Y)}(u_1, a_1) - \mathbb{1}\{\|y\|_\infty \leq Y\} \lambda_j^{h(\bar{y})}(u_1, a_1))^2\right)^{1/2} \\
&\qquad \leq \|\alpha\| \Pr\left(u_1 \in (\ell(y)' a_1, \ell(y)' a_1 + \gamma \iota' a_1]\right) \\
&\qquad \leq \|\alpha\| \mathrm{E}\Big(\Pr\left(u_1 \in (\ell(y)' a_1, \ell(y)' a_1 + \gamma \|\alpha\|_1] \mid a_1\right)\Big) \\
&\qquad \leq \gamma M \|\alpha\|_1^2 \\
&\qquad = \nu.
\end{aligned}
$$

Next, if $e_i'y > Y$ then

$$
\begin{aligned}
\mathrm{E}\left((\lambda_j^{\ell(y\wedge Y)}(u_1,a_1) - \mathbb{1}\{\|y\|_\infty \leq Y\}\lambda_j^{h(\bar{y})}(u_1,a_1))^2\right)^{1/2} \\
&\leq \|\alpha\| \Pr\left(u_1 > a_1'(y \wedge Y)\right) \\
&\leq \|\alpha\| \Pr(u_1 > Y a_1' e_i) \\
&\leq \|\alpha\| \big(\Pr(a_1' e_i \leq A) + \Pr(u_1 > YA) \big) \\
&\leq \nu.
\end{aligned}
$$

Finally, the set $\{(\mathbb{1}\{\|y\|_\infty \leq Y\}\lambda_j^{h(y)}, \lambda_j^{\ell(y\wedge Y)})\}_{y\in\mathbb{R}_+^m}$ has $N_\nu = (\lfloor Y/\gamma \rfloor + 1)^m$ elements. Note that $E(u_1) < \infty$ implies $U < 2\|\alpha\|/\nu$, for all sufficiently small ν. Hence, for small enough ν we have $N_\nu < (\lfloor 2\|\alpha\|/(\gamma\nu A)\rfloor + 1)^m \leq (\lfloor 2\|\alpha\|_1^3 M/(\nu^2 A)\rfloor + 1)^m \leq C/\nu^{2m}$, which implies that $\int_{\nu=0}^1 \sqrt{\log N_\nu} d\nu < \infty$. □

Lemma 6. *For all $\delta > 0$ there exist $\epsilon > 0$ such that* $\mathrm{E}\left(\sup_{y\in B_\epsilon(y_\infty^\beta)} \|\dot{\Lambda}_t^b(y) - \dot{\Lambda}_\infty^b(y) - \dot{\Lambda}_t^b(y_\infty^\beta) + \dot{\Lambda}_\infty^b(y_\infty^\beta)\|^2\right) \leq$ *δ/t for all $t \in \mathbb{N}$.*

Proof. The proof is similar to the proof of Lemma 5, except with

$$\lambda_j^y(u_1,a_1) \equiv (\mathbb{1}\{\Delta_1(y) > 0\} - \mathbb{1}\{\Delta_1(y_\infty^\beta) > 0\})e_j'a_1.$$

Modifying the proof of Lemma 5 establishes that the bracketing integral of $\{\lambda_j^y\}_{y\in B_\epsilon(y_\infty^\beta), j\in[m]}$ is uniformly bounded in $\epsilon \in [0,1]$. Further, λ_j^y is bounded by envelope function $\bar{\lambda}_j^y$, where

$$
\begin{aligned}
\bar{\lambda}_j^y(u_1,a_1) &\equiv (\mathbb{1}\{\Delta_1(\underline{y}) > 0\} - \mathbb{1}\{\Delta_1(\bar{y}) > 0\})e_j'a_1, \\
\underline{y} &\equiv y \wedge y_\infty^\beta, \\
\text{and} \quad \bar{y} &\equiv y \vee y_\infty^\beta.
\end{aligned}
$$

I will now show that the second moment of this envelope can be made arbitrarily small. First, Assumption 6 and Lemma 1 yield the following:

$$\begin{aligned}
\mathrm{E}(\bar{\lambda}_j^y(u_1,a_1)^2) &= \mathrm{E}\left((\mathbb{1}\{\Delta_1(\underline{y})>0\} - \mathbb{1}\{\Delta_1(\bar{y})>0\})(e_j'a_1)^2\right) \\
&\leq \|\alpha\| e_j' \,\mathrm{E}\left((\mathbb{1}\{\Delta_1(\underline{y})>0\} - \mathbb{1}\{\Delta_1(\bar{y})>0\})a_1\right) \\
&= \|\alpha\| e_j' \int_{\gamma=0}^{1} \tfrac{\partial}{\partial\gamma}\,\mathrm{E}\left(\mathbb{1}\{\Delta_1(\bar{y}+\gamma(\underline{y}-\bar{y}))>0\}a_1\right)d\gamma \\
&= -\|\alpha\| e_j' \int_{\gamma=0}^{1} \ddot{\Lambda}_\infty(\bar{y}+\gamma(\underline{y}-\bar{y}))(\underline{y}-\bar{y})d\gamma \\
&= -\|\alpha\| e_j' \ddot{\Lambda}_\infty(\bar{y}+\bar{\gamma}(\underline{y}-\bar{y}))(\underline{y}-\bar{y}),
\end{aligned}$$

for some $\bar{\gamma}\in[0,1]$. And since we constrain $y\in B_\epsilon(y_\infty^\beta)$, it follows that $\ddot{\Lambda}_\infty(\bar{y}+\bar{\gamma}(\underline{y}-\bar{y}))\to\ddot{\Lambda}_\infty(y_\infty^\beta)$ and $(\underline{y}-\bar{y})\to 0$ as $\epsilon\to 0$. Accordingly, $\mathrm{E}(\bar{\lambda}_j^y(u_1,a_1)^2)\to 0$ as $\epsilon\to 0$, which with theorems 2.14.2 and 2.14.5 of van der Vaart and Wellner (1996) establishes the result. □

Lemma 7. *For any compact $\Omega\subset\mathbb{R}_+^m$, there exists $C>0$ such that*
$\mathrm{E}\left(\sup_{y,\bar{y}\in\Omega}|\Lambda_t^b(y)-\Lambda_\infty^b(y)-\Lambda_t^{\bar{b}}(\bar{y})+\Lambda_\infty^{\bar{b}}(\bar{y})|^2\right)\leq C/t$ *for all $t\in\mathbb{N}$, $b,\bar{b}\in\mathbb{R}_+^m$.*

Proof. Like in the proofs of Lemmas 5 and 6, I will use theorems 2.14.2 and 2.14.5 of van der Vaart and Wellner's (1996). As before, I will cast the problem as an empirical process with a new set of functions:

$$\Lambda_t^b(y)-\Lambda_\infty^b(y)-\Lambda_t^{\bar{b}}(\bar{y})+\Lambda_\infty^{\bar{b}}(\bar{y}) = \sum_{s=1}^{t}\lambda_y^{\bar{y}}(u_s,a_s)/t - \mathrm{E}(\lambda_y^{\bar{y}}(u_1,a_1)),$$
$$\text{where}\quad \lambda_y^{\bar{y}}(u_1,a_1)\equiv(u_1-a_1'y)^+-(u_1-a_1'\bar{y})^+.$$

Note that $|\lambda_y^{\bar{y}}(u_1,a_1)|\leq\|\alpha\|\mathrm{diam}(\Omega)<\infty$. Accordingly, it will suffice to show that the bracketing integral of $\{\lambda_j^y\}_{y\in\Omega,\bar{y}\in\Omega}$ is finite.

To bound this bracketing integral, define m-dimensional grid $G\equiv\gamma\mathbb{Z}^m$, where $\gamma\equiv\nu/(4\|\alpha\|\|\iota\|)$. Next, let $\ell(y)\equiv\max\{g\in G\mid g\leq y\}$ represent the largest gridpoint that's weakly less than $y\in\mathbb{R}_+^m$ and let $h(y)\equiv\min\{g\in G\mid g>\ell(y)\}$ represent the smallest gridpoint that's strictly larger than

$\ell(y)$. By design, we have

$$\begin{aligned}\lambda_{\ell(y)}^{h(\bar{y})} - \lambda_{h(y)}^{\ell(\bar{y})} &\leq (u_1 - a_1'(y - \gamma\iota))^+ - (u_1 - a_1'(\bar{y} + \gamma\iota))^+ \\ &\quad - (u_1 - a_1'(y + \gamma\iota))^+ + (u_1 - a_1'(\bar{y} - \gamma\iota))^+ \\ &\leq 4\gamma\|\alpha\|\|\iota\| \\ &= \nu.\end{aligned}$$

Accordingly, the pair $(\lambda_{h(y)}^{\ell(\bar{y})}, \lambda_{\ell(y)}^{h(\bar{y})})$ is a ν-bracket that contains $\lambda_y^{\bar{y}}$. Finally, there are only $O(\nu^{2m})$ such brackets; hence, the bracketing integral is finite. □

Lemma 8. *For all sufficiently small $\epsilon > 0$, the random mapping $(j, y) \mapsto \sqrt{t}e_j'(\dot{\Lambda}_t^b(y) - \dot{\Lambda}_\infty^b(y))$, with $j \in [m]$ and $y \in R_+^m$, weakly converges, as $t \to \infty$, to a mean-zero Gaussian process with domain $[m] \times R_+^m$ and covariance function $\Sigma_{j\bar{j}}(y, \bar{y}) \equiv e_j' \operatorname{E}(\mathbb{1}\{\Delta_1(y) > 0\}\mathbb{1}\{\Delta_1(\bar{y}) > 0\}a_1 a_1') - \operatorname{E}(\mathbb{1}\{\Delta_1(y) > 0\}a_1) \operatorname{E}(\mathbb{1}\{\Delta_1(\bar{y}) > 0\}a_1')e_{\bar{j}}$.*

Proof. This is a direct application of theorem 2.3 of Kosorok, a classical empirical processes result. The proof of Lemma 5 establishes that the corresponding bracketing integral is finite. □