Online Appendix: Random Evolving Lotteries and Intrinsic Preference for Information[†]

Faruk Gul

Paulo Natenzon

and

Wolfgang Pesendorfer[‡]

March 2021

Abstract

In this online appendix, we extend our formal framework to study sequential choice and we allow agents to take an action in every period. We provide a recursive formulation of our model and show that choices are dynamically consistent. We include an application of this recursive setting to address the "ostrich effect."

[†] This research was supported by grant SES-1426252 from the National Science Foundation. We thank Yusufcan Masatlioglu, Erkut Ozbay, the editor, and three anonymous referees for insightful comments and suggestions.

[‡] Faruk Gul, Department of Economics, Princeton University, fgul@princeton.edu; Paulo Natenzon, Olin Business School, Washington University in St. Louis, pnatenzon@wustl.edu; Wolfgang Pesendorfer, Department of Economics, Princeton University, pesendor@princeton.edu.

1. Recursive Choice

In the main body of the paper, we have assumed that the agent chooses an REL at time 0 and makes no further decisions in subsequent periods. In this appendix, we consider a sequential choice setting to analyze the dynamic behavior of PTU agents. In a recursive decision problem, the agent takes an action every period. We let $\mathcal{B}_t = \{1, \ldots, n_t\}$ denote the set of feasible period-t actions and let $\mathcal{B} = \mathcal{B}_1 \times \cdots \times \mathcal{B}_N$.

At the beginning of each period, the agent chooses an action b_t and a random state $\omega_t \in \Omega_t$ occurs. Again, we let $\Omega := \Omega_1 \times \cdots \times \Omega_N$. The action b_1 yields a probability over states in period 1. In every subsequent period t, the action b_t and the state ω_{t-1}, realized in the previous period, determine the probability $\mathcal{Q}_{b_t}(\cdot \,|\, \omega_{t-1})$ over states in period t. We assume that every $\mathcal{Q}_{b_t}(\cdot \,|\, \omega_{t-1})$ has a finite support. Each ω_t contains all of the payoff relevant information revealed during the first t periods. Hence, we refer to the states also as histories. Any history $\omega_t \in \Omega_t$ yields a lottery $\phi_t(\omega_t) \in L$ in period t and, therefore, any $\omega_N \in \Omega_N$ also yields a path $\phi(\omega_N) = \phi(\omega) = (\phi_1(\omega_1), \ldots, \phi_N(\omega_N))$. We let $\phi_{2t}(\omega_t)$ denote the second coordinate, the prize lottery β, of $\phi_t(\omega_t) = (\alpha, \beta)$.

A *strategy* is a collection $\sigma = (\sigma_1, \ldots, \sigma_N) \in \Sigma$ such that $\sigma_1 \in \mathcal{B}_1$ and $\sigma_t : \Omega_t \to \mathcal{B}_t$. We write $\mathcal{Q}_{\sigma_{t+1}}(\cdot \,|\, \omega_t)$ rather than the more cumbersome $\mathcal{Q}_{\sigma_{t+1}(\omega_t)}(\cdot \,|\, \omega_t)$. Clearly, any strategy σ induces a probability Q_σ over evolving lotteries: let

$$\mathcal{Q}_\sigma(\omega) = \mathcal{Q}_{\sigma_1}(\omega_1) \times \cdots \times \mathcal{Q}_{\sigma_{t+1}}(\omega_{t+1} \,|\, \omega_t) \times \cdots \times \mathcal{Q}_{\sigma_{N-1}}(\omega_N \,|\, \omega_{N-1})$$

Then Q_σ, the probability that σ induces on evolving lotteries, is defined as follows:

$$Q_\sigma(x) = \mathcal{Q}_\sigma \phi^{-1}(x)$$

Let $\mathcal{Q}_\sigma(\omega_t) = \mathcal{Q}_\sigma\{\hat{\omega} : \hat{\omega}_t = \omega_t\}$. To ensure that Q_σ is a REL, we impose the martingale property: for all $t < N$, $b_{t+1} \in \mathcal{B}_{t+1}$ and $\omega_t \in \Omega_t$ such that $\mathcal{Q}_\sigma(\omega_t) > 0$,

$$\sum_{(\alpha,\beta) \in L} \beta \mathcal{Q}_{b_{t+1}}(\phi_{t+1}^{-1}(\alpha, \beta) \,|\, \omega_t) = \phi_{t2}(\omega_t)$$

Let $\mathcal{Q}$ denote the entire collection of state probabilities. Then, a *recursive risk consumption problem* (RRCP) is a collection $\mathcal{D} = \{\Omega, \mathcal{B}, \mathcal{Q}, \phi\}$. Below, we describe the optimal solution to an RRCP for a PTU $(u, v, \lambda, \theta_h, \theta_\ell)$. Recall that λ is additive and, therefore, the weights λ_t it assigns to each individual period t fully characterize it. Define:

$$W_N(\omega_{N-1}) = \max_{b_N \in \mathcal{B}_N} \sum_{\omega_N} \left[\lambda_N(1-\theta_h-\theta_\ell)vu\phi_N(\omega_N) + \theta_h\overline{vu}\phi(\omega_N) + \theta_\ell\underline{vu}\phi(\omega_N)\right] \mathcal{Q}_{b_N}(\omega_N|\omega_{N-1})$$

Hence, W_N is the maximal payoff the PTU agent can enjoy in period N, given that the state ω_{t-1} occurred in period $t-1$. Then, for $1 < t < N$, the corresponding maximal payoff is:

$$W_t(\omega_{t-1}) = \max_{b_t \in \mathcal{B}_t} \sum_{\omega_t} \left[\lambda_t(1-\theta_h-\theta_\ell)vu\phi_t(\omega_t) + W_{t+1}(\omega_t)\right] \mathcal{Q}_{b_t}(\omega_t \mid \omega_{t-1})$$

Finally, the maximal period 1 payoff of the agent given the is:

$$W_1 = \max_{b_1 \in \mathcal{B}_1} \sum_{\omega_1} \left[\lambda_1(1-\theta_h-\theta_\ell)vu\phi_1(\omega_1) + W_2(\omega_1)\right] \mathcal{Q}_{b_1}(\omega_1)$$

A strategy σ is sequentially optimal if $\sigma_{t+1}(\omega_t)$ solves the maximization problem that defines $W_{t+1}(\omega_t)$ for all t, ω_t and σ_1 solves the maximization problem that defines W_1.

Proposition A1 below establishes dynamic consistency; that is, the agent's sequentially optimal choice is also ex ante optimal. Its proof is straightforward and omitted.

Proposition A1: *For any PTU V and RRCP $\mathcal{D}$, $W_1 = \max_{\sigma \in \Sigma} V(Q_\sigma)$.*

The proposition above shows that optimal sequential choice is equivalent to choosing an optimal REL, ex ante, from the set of RELs that can be generated from the given decision problem. Thus, our definition of W is consistent with the agent's ex ante preference and yields dynamically consistent behavior.

1.1 The Ostrich Effect

Agents exhibit the Ostrich effect if they seek information after good news and reject information after bad news. The simplest example to capture this effect with four periods and two signals is contained in the main text. Here, we define a version of that simple example for arbitrary $N \geq 4$ and show that the ostrich effect holds generally. As in the simple example, in each period $i = 2, \ldots, N-1$, the agent decides whether to observe signal i. The signals are identically distributed and conditionally independent given the outcome. The initial (i.e., period 1) probability of outcome 1 is $a \in (0, 1)$ and, conditional on outcome $i \in \{0, 1\}$, the probability of getting signal i is $d > 1/2$. Each period, the agent has two choices: she can acquire information ($b_t = 1$) or not ($b_t = 0$). We let $\omega_t = (j_1, \ldots, j_t)$ be the t-period history of information outcomes where $j_s = 1$ (-1) if the news in period s is good (bad) and $j_s = 0$ if the agent does not acquire information in period s. Hence, $j_1 = 0$ for all histories ω. The probability of the good prize in period t given the history ω_t is $\phi(\omega_t)$.

We say that the agent displays the ostrich effect if there is some sequentially optimal strategy σ and a history ω_{t-1} that occurs with positive probability given σ such that σ gets information in period t after history ω_{t-1} and, in period $t+1$, gets information if and only if there is good news in period t.

Proposition A2: *Let $(u, v, \lambda, \theta_h, \theta_\ell)$ be a PTU with linear v and $\theta_h, \theta_\ell > 0$. (i) for all $\omega_t = (0, j_1, \ldots, j_{t-1}, 1)$, there is a sequentially optimal σ such that $\sigma_{t+1}(\omega_t) = 1$. (ii) Either it is optimal to never get information or the agent displays the ostrich effect.*

Proposition A2 establishes that it is always optimal to get more information after good news. It also shows that in some contingencies, the agent must get additional information after good news and avoid additional information after bad news. If the decision maker receives good news in period t, then there is no chance that additional information in period $t+1$ can create a new low and after some histories may create a new high. Conversely, if the news in period t is bad, the new information might create a new low but cannot create a new high. After certain balanced histories; that is, histories in which the agent has received roughly the same amount of good and bad news, the desire to create new highs and avoid new lows ensures that the agent displays the ostrich effect.

The proposition above allows for the possibility that the agent may never wish to get information. However, fixing all other parameters, increasing θ_h or decreasing θ_ℓ or decreasing a sufficiently will rule out this possibility and ensure that the agent displays the ostrich effect. Finally, increasing N increases the option value of information and hence increases the range of the other parameters for which the ostrich effect is guaranteed.

Proposition A3 below, identifies another behavioral pattern of PTU agents who savor peaks and dread troughs. Consider again the above information acquisition problem but with more general signals. Specifically, the signals are i.i.d. conditional on the true value of the prize and have finitely many possible realizations. The agent must decide whether to acquire information and get the updated prize lottery (recall that the consumption lottery is fixed and constant regardless of the agent's decisions) or stay uninformed and get the same prize lottery as in the previous period.

As above, let ω_t be the t-period history of information outcomes and let $\phi(\omega_t)$ be the probability of the good outcome given ω_t. Let $\pi(\omega_t) = \max\{v(u(\phi(\omega_s)) \,|\, s \le t\}$ and $\tau(\omega_t) = \min\{v(u(\phi(\omega_s)) \,|\, s \le t\}$. Proposition A3 shows that for any fixed ϕ_t, histories with lower peaks and lower troughs yield greater willingness to acquire information:

Proposition A3: *Let $(u, v, \lambda, \theta_h, \theta_\ell)$ be a PTU with $\theta_h, \theta_\ell > 0$. Suppose $\phi_t(\omega_t) = \phi_t(\hat{\omega}_t)$, $\pi(\omega_t) \ge \pi(\hat{\omega}_t)$ and $\tau(\omega_t) \ge \tau(\hat{\omega}_t)$. Then, if it is optimal to get information after history ω_t, it is also optimal to get information after history $\hat{\omega}_t$.*

Proposition A3 holds fixed the choice problem but varies the peak-trough history. It shows that if the utility of the current prize lottery is closer to the past peak and further from past trough experiences, the agent is more inclined to choose the informative signal. Holding the current prize lottery fixed neutralizes the effect of the curvature of v and isolates the effect of past experience on the agent's choice behavior.

To provide intuition for Proposition A3, define the period-t utility flow ι_t as follows:

$$\begin{aligned}\iota_t(\omega_t) = {} & (1-\theta_h-\theta_\ell)\lambda_t v(\phi(\omega_t)) \\ & + \theta_h(v(\pi(\omega_t)) - v(\pi(\omega_{t-1})) + \theta_\ell(v(\tau(\omega_t)) - v(\tau(\omega_{t-1}))\end{aligned}$$

The first term on the right hand side is the flow utility of the period-t prize lottery. The second and third terms are the utilities associated with reaching a new peak or a new trough in period t; their sum is positive if a new peak is reached, negative if a new trough is reached and zero otherwise. In the final period, period N, the agent has no choice; let $W_N = E[\iota_N]$ where the expectation is taken over the lottery that reveals all remaining uncertainty given the state ω_{N-1}. In preceding periods, the agent chooses the informative signal if the expectation of $\iota_t + W_{t+1}$ is greater under the informative than under the uninformative signal. If we ignore the continuation, the result in Proposition A3 is straightforward: if peak and trough are smaller, then the instantaneous payoff for the informative signal is higher since the chance that the prize lottery reaches a new peak in period t is larger and the chance it reaches a new trough is smaller. For the uninformative signal, by contrast, the instantaneous payoff remains unaffected by changes in historical peaks and troughs. Thus, if we only consider the instantaneous utility, lower peaks and troughs tilt the trade-off in favor of the informative signal. Proposition A2 shows that this conclusion continues to hold even if we take the continuation into account and even if all information is revealed in the final period.

2. Proof of Proposition A2

Let $P = Q_\sigma$ and $\hat{P} = Q_{\hat{\sigma}}$ for strategies σ and $\hat{\sigma}$. Throughout the proof of part (ii), we assume that it is not optimal to never get information.

For any $\omega_t = (0, j_2, \ldots, j_{t-1})$, let σ^{ω_t} be the strategy such that the agent does in period $s \geq t$ after history $\omega_s = (0, j_2, \ldots, j_{t-1}, \ldots, j_{s-1})$ what she would have done after history $(0, 0, j_2, \ldots, j_{t-1}, \ldots, j_{s-1})$ with strategy σ for all such ω_{s-1} and gets no information after any history $\omega_{N-2} = (0, j_2, \ldots, j_{N-2})$. After any history that does not follow ω_t, σ^{ω_t} chooses the same action as σ.

For any $\omega_t = (0, j_2, \ldots, j_t)$ and $s \leq t$, let $\mu_s(\omega_t) = \sum_{\hat{s}=0}^{s} j_{\hat{s}}$, $\overline{\mu}(\omega_t) = \max_{s \leq t} \mu_s(\omega_t)$ and $\underline{\mu}(\omega_t) = \min_{s \leq t} \mu_s(\omega_t)$.

Step 1: If σ is optimal, then $\sigma_2(0) = 1$.

Let σ be an optimal strategy. By assumption, there is some t such that $\sigma_{t+1}(0, \ldots, 0) = 1$. Consider the strategy σ^{ω_1} for $\omega_1 = (0)$. Hence σ^{ω_1} is also optimal. Next, take any history ω_{N-2} such that $\mu_s(\omega_{N-2}) \geq 0$ for all $s \leq N-2$ and $\mu_{N-2}(\omega_{N-2}) > 0$. Since σ^{ω_1} gets information in period $t < N-2$, there must be some such a history that occurs with positive probability given σ^{ω_1}. Let $\hat{\sigma}$ be the strategy derived from σ^{ω_1} by replacing $\sigma_{N-1}(\omega_{N-2})$ for that history ω_{N-2} with $\hat{\sigma}_{N-1}(\omega_{N-2}) = 1$. Clearly, $\hat{P}_\ell = P_\ell$ and $\hat{P}_h$ stochastically dominates P_h, contradicting the optimality of σ.

Step 2: If σ is sequentially optimal and $\underline{\mu}(\omega_t) < \mu_t(\omega_t) = \overline{\mu}(\omega_t)$, then $\sigma_{t+1}(\omega_t) = 1$.

The proof is similar to the proof of Step 1. Assume that σ is sequentially optimal. By Step 2, $\sigma_2 = 1$. Then, take any history ω_{N-2} such that $\mu_s(\omega_{N-2}) \geq 0$ for all $s \leq N-2$ and $\mu_{N-2}(\omega_{N-2}) > 0$. Since $\sigma_2 = 1$, such an ω_{N-2} occurs with positive probability given σ^{ω_t}. Then, consider the strategy $\hat{\sigma}$ that is derived from σ^{ω_t} by replacing $\sigma_{N-1}(\omega_{N-2}) = 0$ with $\hat{\sigma}_{N-1}(\omega_{N-2}) = 1$. Clearly, $\hat{P}_\ell = P_\ell$ and $\hat{P}_h$ stochastically dominates P_h, contradicting the sequential optimality of σ.

For any $k \geq 2$, odd let $\omega^o_{2k-1} = (0, 1, -1, \ldots, 1, -1)$.

Step 3: If, for some $k \geq 2$, there is a sequentially optimal σ such that $\sigma_{2k}(\omega^o_{2k-1}) = 1$, then $\sigma_{2k-2}(\omega^o_{2k-3}) = 1$ for every sequentially optimal σ.

Let σ be any sequentially optimal strategy such that $\sigma_{2k-2}(\omega^o_{2k-3}) = 0$. Suppose that $\sigma_{2k-1}(\omega_{2k-2}) = 1$ for $\omega_{2k-2} = (0, 1, -1, \ldots, 1, -1, 0)$. Then, the proof of Step 1 establishes that $\sigma^{\omega^o_{2k-3}}$ yields the same payoff as σ after every history and is therefore also sequentially optimal. Moreover, $\sigma^{\omega^o_{2k-3}}$ gets no information after the $N-2$-period history that follows ω^o_{2k-2} in which there is good news in every period $t > 2k-3$, contradicting Step 2. Hence, $\omega_{2k-2} = (0, 1, -1, \ldots, 1, -1, 0) = 0$. Repeating the last argument establishes that $\sigma_{2k}(0, 1, -1, \ldots, 1, -1, 0, 0) = 0$ for every sequentially optimal σ which implies that $\sigma_{2k}(0, 1, -1, \ldots, 1, -1, 1, -1) = 0$ for every sequentially optimal σ, a contradiction.

If $N = 4$, then the decision maker displays the ostrich effect. To see this, let $N = 4$ and consider any sequentially optimal σ. Then, by Step 1, $\sigma_2(0) = 1$. By Step 2, $\sigma_3(0, 1) = 1$. We claim that $\sigma_3(0, -1) = 0$. To see why this is the case, assume σ always get information while $\hat{\sigma}$ gets information after every history except $(0, -1)$. Then, it is easy to see that P_ℓ is stochastically dominated by $\hat{P}_\ell$ and $\hat{P}_h = P_h$. Hence, $V(\hat{P}) > V(P)$ contradicting the optimality of σ. That is; $\sigma_2 = \sigma_3(0, 1) = 1$ and $\sigma_3(0, 1) = 0$ and therefore, the agent displays the ostrich effect.

Henceforth, we will assume that $N > 4$. Define,

$$\mathcal{K} = \{k \geq 2 : Q_\sigma(\omega^o_{2k-1}) > 0 \text{ and } \sigma_{2k}(\omega^o_{2k-1}) = 1 \text{ for some optimal } \sigma\}$$

First, assume that $\mathcal{K} = \emptyset$. Suppose σ is any optimal strategy. Then, by Step 2, $\sigma_2 = 1$ and therefore $\omega_1 = (0, 1)$ occurs with probability. By Step 3, $\sigma_4(0, 1, 1) = 1$ and since $2 \notin \mathcal{K}$, $\hat{\sigma}_4(0, 1, -1) = 0$. Hence, the decision maker displays the ostrich effect.

Next, assume $\mathcal{K} \neq \emptyset$ and let $\kappa = \max \mathcal{K}$. We note that N must be greater or equal to $2\kappa + 2$. If not; that is, if $N = 2\kappa + 1$, then let σ be any optimal strategy such that ω^o_{2k-1} occurs with positive probability given σ and $\sigma_{2k}(\omega^o_{2k-1}) = \sigma_{N-1}(\omega^o_{2k-1}) = 1$ consider the strategy $\hat{\sigma}$ such that $\hat{\sigma}_{t+1}(\omega_t) = \sigma_{t+1}(\omega_t)$ for all $t \neq 2\kappa$ or ($t = 2\kappa$ and $\omega_{t-1} \neq \omega^o_{t-1}$) and let $\hat{\sigma}_{2\kappa}(\omega^o_{2\kappa-1}) = 0$. Note that $\hat{P}_h = P_h$ and $\hat{P}_\ell$ stochastically dominates P_ℓ and therefore $V(\hat{P}) > V(P)$, contradicting the optimality of σ and proving that $N \geq 2\kappa + 2$.

Suppose $N = 2\kappa + 2$. Let σ be an optimal strategy such that $\omega^o_{2\kappa-1}$ occurs with positive probability given σ and $\sigma_{2\kappa}(\omega^o_{2\kappa-1}) = 1$. Then, if the news in period 2κ is good, it is optimal for the agent to get information in period $2\kappa+1$ by Step 2. Conversely, if the news in period 2κ is bad, it is optimal for the agent not to get information in period $2\kappa + 1$ since the additional information cannot increase $\overline{vu}(x)$ but may decrease $\underline{vu}(x)$. Hence, the agent displays the ostrich effect.

If $N > 2\kappa + 2$, then by the definition of κ, there is an optimal strategy σ such that $\omega^o_{2\kappa-1}$ occurs with positive probability given σ and $\sigma_{2\kappa}(\omega^o_{2\kappa-1}) = 1$. Then, history $\omega_{2\kappa} = (0, 1, -1, \ldots, 1, -1, 1)$ also occurs with positive probability given σ. By Step 3, $\sigma_{2\kappa+1}(\omega_{2\kappa}) = 1$ and $\hat{\sigma}_{2\kappa+2}(0, 1, -1, \ldots, 1, -1, 1, 1) = 1$. Hence, $\omega^o_{2\kappa+1}$ also occurs with positive probability given σ and therefore, by the definition of κ, $\hat{\sigma}(\omega^o_{2\kappa+1}) = 0$. Hence, the decision maker displays the ostrich effect. This concludes the proof of part (ii) of the proposition.

To prove part (i) of the proposition, take any $\omega_t = (0, j_2, \ldots, j_{t-1}, 1)$ and assume there is no sequentially optimal strategy σ such that $\sigma_{t+1}(\omega_t) = 1$. Let σ be any sequentially optimal strategy. (Hence, $\sigma_{t+1}(\omega_t) = 0$.) Suppose there some history $\omega_s = (0, j_2, \ldots, j_{t-1}, 1, 0, \ldots, 0)$ such that $\sigma_{s+1}(\omega_s) = 1$. Then, let $\sigma^1 = \sigma^{\omega_t}$, $\sigma^2 = \sigma^{1\omega_t}$, $\sigma^3 = \sigma^{2\omega_t}$ and so on.

Note that $\sigma^{s-t}_{s+1}(\omega_t) = 1$. We noted in the proof of Step 1 that σ^{ω_t} yields the same distribution of peaks and troughs as σ and therefore, results in the same utility. Hence, σ^{s-t} yields the same utility as σ, and therefore, is also sequentially optimal, contradicting the fact that there is no sequentially optimal strategy σ such that $\sigma_{t+1}(\omega_t) = 1$. $\square$

3. Proof of Proposition A3

Assume that it is optimal to get information after history ω_t but not optimal to get information after history $\hat{\omega}_t$ where $\pi(\omega_t) \geq \pi(\hat{\omega}_t)$, $\tau(\omega_t) \geq \tau(\hat{\omega}_t)$ and $\phi_t(\omega_t) = \phi_t(\hat{\omega}_t)$.

First, we claim that getting no information in any period following $\hat{\omega}_t$ is optimal. To see why, suppose that there is a sequentially optimal strategy, $\hat{\sigma}$, such that $\hat{\sigma}_{t+1}(\hat{\omega}_t) = 0$ and $\hat{\sigma}_{s+1}(\hat{\omega}_{\hat{s}}) = 1$ for some $\hat{s} > t$. Let s be the smallest such $\hat{s}$. Then, repeating the construction in Step 1 of the proof of Proposition 6 yields a strategy $\sigma^1 = \hat{\sigma}^{\omega_t}$ that yields the same payoff as $\hat{\sigma}$ such that $\sigma^1_{t+1}(\hat{\omega}_t) = 0$ and σ^1 gets no information in period $\hat{s} = t+1, \ldots, s-1$ following history $\hat{\omega}_t$ and gets information in period s. Repeating this construction with $\hat{\sigma}^{\omega_t}$ replacing $\hat{\sigma}$ yields an optimal strategy, σ^2 that does not get information after history ω_t until period $s-1$ and gets information in period $s-1$. Hence, repeating this process again and again yields a strategy σ^{s-t} that is optimal and gets information in period $t+1$ after history ω_t, contradicting the assumption that no such strategy exists.

Since it is optimal to get information after ω_t, there must exist a sequentially optimal strategy σ that yields a weakly higher payoff than never getting information after ω_t such that $\sigma_{t+1}(\omega_t) = 1$. Let $P = Q_\sigma$.

$$\theta_h[E_{P_h}v - \pi(\omega_t)] + \theta_\ell[E_{P_\ell}v - \tau(\omega_t)] \geq 0$$

Let $\pi(\omega_t)^+$ denote the set of all $\beta \in [0,1]$ such that $P_h(\beta) > 0$ and $v(\beta) > \pi(\omega_t)$ and let $\tau(\omega_t)^-$ be the set of all $\beta \in [0,1]$ such that $P_\ell(\beta) > 0$ and $v(\beta) < \pi(\omega_t)$. Then, since P_h never delivers an outcome below $\pi(\omega_t)$ and P_ℓ never delivers an outcome above $\tau(\omega_t)$, the equation above is equivalent to

$$\theta_h \sum_{\pi(\omega_t)^+} [v(\beta) - \pi(\omega_t)]P_h(\beta) + \theta_\ell \sum_{\tau(\omega_t)^-} [v(\beta) - \tau(\omega_t)]P_\ell(\beta) \geq 0 \qquad (*)$$

Now, assume that the decision maker behaves after history $\hat{\omega}_t$ as she did with strategy σ after history ω_t; that is, the behavior assumed in equation $(*)$. This strategy is not optimal since it entails getting information in period $t+1$. Furthermore, as we have noted above, never getting information after history $\hat{\omega}_t$ is optimal. Hence, arguing as we did in deriving equation $(*)$ above yields

$$\theta_h \sum_{\pi(\hat{\omega}_t)^+} [v(\beta) - \pi(\hat{\omega}_t)]P_h(\beta) + \theta_\ell \sum_{\tau(\hat{\omega}_t)^-} [v(\beta) - \tau(\hat{\omega}_t)]P_\ell(\beta) < 0 \qquad (**)$$

Equations $(*)$ and $(**)$ yield

$$\begin{aligned} \theta_h \sum_{\pi(\omega_t)^+} [\pi(\omega_t) - \pi(\hat{\omega}_t)]P_h(\beta) + \theta_h \sum_{\pi(\hat{\omega}_t)^+ \setminus \pi(\omega_t)^+} [v(\beta) - \pi(\hat{\omega}_t)]P_h(\beta) \\ + \theta_\ell \sum_{\tau(\omega_t)^- \setminus \tau(\hat{\omega}_t)^-} [\tau(\omega_t) - v(\beta)]P_\ell(\beta) + \theta_\ell \sum_{\tau(\hat{\omega}_t)^-} [\tau(\omega_t) - \tau(\hat{\omega}_t)]P_h(\beta) < 0 \end{aligned}$$

The last inequality establishes a contradiction since all of the summands above are nonnegative. $\square$