Optimal Stopping with Randomly Arriving Opportunities to Stop

Josha A. Dekker[1] Roger J.A. Laeven[1]

John G.M. Schoenmakers[2] Michel H. Vellekoop[1]

[1]University of Amsterdam, Actuarial Science and Mathematical Finance group

[2]Weierstrass Institute for Applied Analysis and Stochastics

April 5, 2024

University of Amsterdam
Amsterdam School of Economics

Classical optimal stopping problems

- Optimally stop a reward process Z_t at some predetermined moments:
  - Bermudan, stop rewards Z_t at predetermined $t \in \{t_1, ..., t_k\}$.
  - American, stop rewards Z_t at any $t \in [0, T]$.
- Moments may not be predetermined in practice:
  - Practical constraints: *illiquid markets, real options*.
  - Inherent randomness: *catastrophe derivatives, uncertainty of success*.
- We look at a generalization of Bermudan stopping problems:
  - (Some of) the opportunities arrive randomly
  - The terms opportunities, buyers and arrivals will be associated with each-other for the purpose of this talk.

Leading example: valuable asset

- Suppose you have an exclusive, illiquidly traded house or piece of art.
- May be sold at the prevailing market price X_t.
- Selling requires a buyer to arrive.
- Attach a personal, non-monetary value C, e.g. derived from utility, so only sell if $X_t > C$.
- Reward process: $Z_t = e^{-rt}(X_t - C)^+$, a call-option arises naturally here.
- Total value of the house: value of optimally stopped reward process $+ C$.

Feedback

- Arrivals of opportunities and prices may be interrelated.
- E.g. in our example
  - If prices are high, fewer buyers may arrive.
  - When buyers arrive this may reveal new information, affecting prices.

Literature

Special case: opportunities arrive according to a Poisson process.

- Dupuis and Wang (2002): Perpetual call option, independent arrivals
- Hobson (2021): Shape of the value function, arrival rate depends on price dynamics
- Lange et al. (2020): develop a finite difference method for independent arrivals.

Related problems:

- Matsumoto (2006): Optimal portfolio selection, independent randomly arriving adjustment dates
- Pham and Tankov (2008, 2009): Optimal investment and consumption choice

Our contribution:

- Tackle more general feedback/dependence structures

Problem formulation

- Sequence of $(\mathcal{F}_t)_{t\geq 0}$-stopping times $0 < \tau_1 < \tau_2 < \cdots$, e.g. event times of a Poisson process, where $\tau_k \to \infty$ a.s.
- Fixed time-horizon $T < \infty$.
- $Z_t = f(t, X_t)$ non-negative reward process where X_t is a $(\mathcal{F}_t)_{t\geq 0}$-adapted Markov process.
- Problem: optimally stop Z_t at one of the stopping times $\tau_k \leq T$ to receive Z_{τ_k}.
- Question: what is the optimal value $Y_0 = \sup_{\tau \in \mathcal{T}(\tau_1, \tau_2, \dots)} \mathbb{E}[Z_\tau \mathbb{1}_{\tau \leq T}]$?
- Question: what is the corresponding optimal stopping rule?

Numerical approaches

- When we allow for feedback, we are typically thinking about higher-dimensional situations (in the order of dimensions 4-10).
- Some possible approaches would be:
  - Bermudan/American approximation: may be inefficient, especially when there are only few expected arrivals.
  - Equivalent problem: American optimal stopping problem with $\tilde{Z}_t = Z_t \mathbb{1}_{\{\exists k,\ t=\tau_k\}}$: this is a continuous-time optimal stopping problem with làdlàg rewards.
  - Finite-difference or finite-elements methods (such as Lange et al. (2020)): we face the curse of dimensionality, especially in the face of feedback.
- We thus need some other approach, more tailored to the problem at hand!

Changing perspectives

- Observation: The dynamics of interest all happen at the random times.
- We can focus our attention on these random times, jumping from random time to random time.
- New process $\{Z_{\tau_k}\}_{k\in\mathbb{N}}$ and filtration $\{\mathcal{F}_{\tau_k}\}_{k\in\mathbb{N}}$.
- This is now a discrete-time, infinite horizon problem!

Equivalence

- Original problem: $Y_0 = \sup_{\tau \in \mathcal{T}(\tau_1, \tau_2, \ldots)} \mathbb{E}[Z_\tau \mathbb{1}_{\tau \leq T}]$.
- $\mathcal{N}$, the set of discrete $\{\mathcal{F}_{\tau_k}\}_{k \in \mathbb{N}}$-stopping times.

Proposition (Equivalence)

$Y_0 = \sup_{\mathfrak{n} \in \mathcal{N}} \mathbb{E}\left[Z_{\tau_\mathfrak{n}} \mathbb{1}_{\tau_\mathfrak{n} \leq T}\right]$.

Finite-horizon approximation

- Some algorithms require a discrete-time, finite horizon problem
- We can pass to a finite horizon problem:
  - Suppose $\mathbb{E}[\sup_i Z_{\tau_i}^2] \leq B^2 < \infty$
  - For K sufficiently large, $Y_0^{(K)} := \sup_{\tau \in \mathcal{T}(\tau_1,\ldots,\tau_K)} \mathbb{E}[Z_\tau \mathbb{1}_{\tau \leq T}]$ can then be made ε-close to Y_0.
- In fact:

Proposition (Truncated equivalence)

$$Y_0^{(K)} = \sup_{\mathfrak{n} \in \mathcal{N},\ \mathfrak{n} \leq K} \mathbb{E}\left[Z_{\tau_\mathfrak{n}} \mathbb{1}_{\tau_\mathfrak{n} \leq T}\right].$$

Duality

Proposition (Weak and strong duality)

Let $\mathcal{M}_0^{\mathrm{UI}}$ be the collection of uniformly integrable $\{\mathcal{F}_{\tau_k}\}_{k\in\mathbb{N}}$-martingales that start from 0. Suppose that $\mathbb{E}[\sup_i Z_{\tau_i}^2] \leq B^2 < \infty$ and $\sum_{j=0}^{\infty} \mathbb{P}(\tau_j \leq T)^{\alpha} < \infty$ for some $0 < \alpha < 1/2$. Then:

i) *Weak duality: It holds that* $Y_0 = \inf_{M\in\mathcal{M}_0^{\mathrm{UI}}} \mathbb{E}[\sup_i (Z_{\tau_i} - M_i)]$.

ii) *Strong duality: Let* $M_i^{\circ} := \sum_{j=1}^{i} Y_j - \mathbb{E}_{\mathcal{F}_{\tau_{j-1}}}[Y_j]$, $i \geq 1$, $M_0^{\circ} = 0$. *Then it holds that* $M^{\circ} \in \mathcal{M}_0^{\mathrm{UI}}$ *and that* $Y_0 = \sup_{i\geq 0} (Z_{\tau_i} - M_i^{\circ})$ *almost-surely.*

- Interpretation of the new condition: the distribution of $N_T = \sup\{j : \tau_j \leq T\}$ has tails that decay faster than j^{-2}.

Algorithms

We can get analogues of familiar simulation-based algorithms for this random times setting:

- Primal: Longstaff and Schwartz (2001) (Least-Squares Monte-Carlo)
- Primal: Andersen (2000) ("Optimal thresholds")
- Dual: Andersen and Broadie (2004) (requires our duality result and a primal method)

We also show that policy iteration (Kolodko and Schoenmakers, 2006), a recursive method to obtain increasingly better approximate policies:

- extends to the infinite horizon case,
- applies to the random times situation
- and has the necessary convergence properties.

Flexibility of the framework

The above algorithms are examples; other simulation based algorithms could be used, such as:

- local regression approaches
- COS method
- etc.

Model assumptions are minimal, many familiar dynamics can be used, e.g.:

- Underlying: GBM, Ornstein-Uhlenbeck, Heston-SVJ
- Arrivals: doubly stochastic Poisson, (Markovian) Hawkes, generalized intensity
- Interaction: arrivals with threshold, state-dependent probability of success, embedded in a multivariate Hawkes process

Example: Max-call on Poisson random times, set-up

- Benchmark in the literature: Example 1 of Andersen and Broadie (2004):
  - Payoff on two GBMs $(\max\{X_t^1, X_t^2\} - K)^+$, $K = 100$.
  - Usually: Bermudan stopping problem, 9 equidistant stopping opportunities on $[0, T]$ with $T = 3$.
- Random-times counterpart: Poisson arriving stopping opportunities with rate λ
- Extends to feedback

Example: square payoff on Poisson random times, methods

- Primal algorithm: Least-Squares Monte-Carlo
  - Make approximate optimal decisions through estimated continuation values
  - Estimate through backward recursion
- Dual algorithm: Andersen-Broadie dual algorithm, uses our duality result
- Let $X_t^{(1)} = \max\{X_t^1, X_t^2\}$ and $X_t^{(2)} = \min\{X_t^1, X_t^2\}$.
- We then use the regression variables $\phi_i(t)(X_t^{(m)})^j$ and $(X_t^{(1)})^j(X_t^{(2)})^k$, where $i \in \{0,1,2,3,4,5\}$, $j,\ k \in \{0,1,2,3\}$ and $m \in \{1,2\}$ and ϕ_i is the i'th Laguerre polynomial

Example: Max-call on Poisson random times, results

X_0	λ	Primal	(s.e.)	Dual	(s.e.)
90	1	5.6876	0.0076	5.6950	0.0081
90	2	6.8617	0.0080	6.8754	0.0084
90	5	7.6772	0.0083	7.7168	0.0094
100	1	10.5710	0.0101	10.5769	0.0102
100	**2**	**12.3455**	**0.0104**	**12.3607**	**0.0108**
100	5	13.4558	0.0105	13.5195	0.0184
110	1	17.1219	0.0123	17.1269	0.0123
110	2	19.5328	0.0123	19.5600	0.0151
110	5	20.8970	0.0123	20.9724	0.0208

Table 1: $N = 2$, $\mu = r = 5\%$, $\delta = 10\%$, $\sigma = 20\%$ $T = 3$. *Andersen and Broadie (2004) method: 200,000 paths for the regression step in LSMC, 2,000,000 paths to determine the primal price, 1,500 paths to determine the dual price, 10,000 to estimate the dual-martingale along each path. Maximum truncation error is 0.001.*

- Duality gaps typically between 0.05% and 0.5%.

Example: Max-call on Poisson random times, relation to Andersen and Broadie (2004)

- In the random times case for $X_0 = 100$
  - $\lambda = 1$ ($\mathbb{E}[N_T] = 3$): upper-estimate is 10.577
  - $\lambda = 2$ ($\mathbb{E}[N_T] = 6$): upper-estimate is 12.361
  - $\lambda = 5$ ($\mathbb{E}[N_T] = 15$): upper-estimate is 13.520
- In the deterministic case with 9 equidistant opportunities, the lower-estimate is 13.907, which is substantially larger.
- This is a consistent finding
  - Stochasticity of opportunities may have a sizeable effect on the value of the optimal stopping problem.
  - The effect may have either sign, as we show in our forthcoming updated preprint.

Example: Max-call with feedback

- Extend the previous example with feedback from price to arrivals
- Idea: rate $\lambda(t, X_t; \alpha)$, where α determines the degree of feedback
  - $\alpha < 0$: negative feedback, high X_t leads to fewer arrivals
  - $\alpha > 0$: positive feedback, high X_t leads to more arrivals
- "Fair comparison": we want $\mathbb{E}[\lambda(t, X_t; \alpha)] = \lambda$, a constant.
- Use $\lambda(t, x; \alpha) = (1-\alpha)\lambda + 2\alpha F_{X_t}(x)\lambda$ for $\alpha \in [-1, 1]$; $F_{X_t}(x) = \mathbb{P}(X_t \leq x)$.

Example: Max-call with feedback, results

X_0	α	Primal	(s.e.)	Dual	(s.e.)	$\mathbb{E}[\tau^*]$	$\mathbb{E}[\tau_{\text{triv}}]$
90	-1	0.3667	0.0014	0.3667	0.0014	2.8810	2.8699
90	-0.5	2.2571	0.0051	2.2582	0.0051	2.7919	2.6799
90	0	3.0870	0.0058	3.0917	0.0064	2.7537	2.5558
90	0.5	3.4832	0.0060	3.4857	0.0061	2.7441	2.4661
90	1	3.7199	0.0061	3.7244	0.0062	2.7400	2.4008
110	-1	5.4504	0.0054	5.4505	0.0054	1.9068	1.7512
110	-0.5	9.0379	0.0091	9.0389	0.0091	1.9476	1.5867
110	0	10.6375	0.0098	10.6420	0.0099	1.9879	1.4841
110	0.5	11.4425	0.0101	11.4479	0.0101	2.0201	1.4163
110	1	11.8970	0.0101	11.9028	0.0102	2.0205	1.3687

Table 2: $N = 1$, $\mu = r = 5\%$, $\delta = 10\%$, $\sigma = 20\%$ $T = 3$.

- Prices increase with α (here).
- Effect on $\mathbb{E}[\tau^*]$ indeterminate.

Conclusions and future research

- The situation studied extends the case with Bermudan opportunities.
- Allows for general feedback structures, no strong structural conditions: has potential for many financial/economic questions.
- Numerical methods work and show a sizeable effect of stochastic opportunities.
- Preprint: `https://arxiv.org/abs/2311.11098`.
- Currently: studying the effect of stochastic structure on optimal stopping decisions (next iteration of preprint):
  - Effect of stochastic arrivals vs deterministic arrivals.
  - Effect of clustered arrivals.
  - Effect of feedback in either direction between arrivals and the reward process.

ANDERSEN, L. (2000): "A Simple Approach to the Pricing of Bermudan Swaptions in the Multifactor LIBOR Market Model," *Journal of Computational Finance*, 3, 5–32.

ANDERSEN, L. AND M. BROADIE (2004): "Primal-Dual Simulation Algorithm for Pricing Multidimensional American Options," *Management Science*, 50, 1222–1234.

DUPUIS, P. AND H. WANG (2002): "Optimal Stopping with Random Intervention Times," *Advances in Applied Probability*, 34, 141–157.

HOBSON, D. (2021): "The Shape of the Value Function under Poisson Optimal Stopping," *Stochastic Processes and their Applications*, 133, 229–246.

KOLODKO, A. AND J. SCHOENMAKERS (2006): "Iterative Construction of the Optimal Bermudan Stopping Time," *Finance and Stochastics*, 10, 27–49.

LANGE, R. J., D. RALPH, AND K. STORE (2020): "Real-Option Valuation in Multiple Dimensions Using Poisson Optional Stopping Times," *Journal of Financial and Quantitative Analysis*, 55, 653–677.

Longstaff, F. A. and E. S. Schwartz (2001): "Valuing American Options by Simulation: A Simple Least-Squares Approach," *The Review of Financial Studies*, 14, 113–147.

Matsumoto, K. (2006): "Optimal Portfolio of Low Liquid Assets with a Log-Utility Function," *Finance and Stochastics*, 10, 121–145.

Pham, H. and P. Tankov (2008): "A Model of Optimal Consumption under Liquidity Risk with Random Trading Times," *Mathematical Finance*, 18, 613–627.

——— (2009): "A Coupled System of Integrodifferential Equations Arising in Liquidity Risk Model," *Applied Mathematics and Optimization*, 59, 147–173.