Optimal Switching Games for Emissions Trading

Mike Ludkovski
Department of Statistics & Applied Probability
University of California Santa Barbara

IPAM, January 8, 2010

Outline

1. Cap-and-Trade: Producer Perspective
2. Switching Games
3. Correlated Equilibria in CO_2 markets
4. Numerical Illustrations
5. Open Problems

Emissions Trading

- Major new initiatives are underway to introduce CO_2 cap-and-trade schemes that will create new commodity markets.
- AB32 proposal in California; various federal proposals.
- The estimated size of the market is in the hundreds of billions or even trillions of dollars.
- Key regulatory details are still unresolved and undergo active public debate.
- Crucial to understand the financial implications of these initiatives on energy producers.

A New Commodity Market

Compared to existing markets, cap-and-trade is fundamentally different:

- A finite resource is initially allocated and subject to exhaustion.
- A well-defined horizon (e.g. 1 year) exists for each allocation.
- The permit prices converge to deterministic values as horizon approaches.
- Price formation is driven by participant strategies: must be endogenous to any model.
- Game-theoretic aspects emerge in the emissions market.

A New Commodity Market

Compared to existing markets, cap-and-trade is fundamentally different:

- A finite resource is initially allocated and subject to exhaustion.
- A well-defined horizon (e.g. 1 year) exists for each allocation.
- The permit prices converge to deterministic values as horizon approaches.
- Price formation is driven by participant strategies: must be endogenous to any model.
- Game-theoretic aspects emerge in the emissions market.

Effect on Producers

- The foremost constituency affected by cap-and-trade would be energy producers.
- The net profit of energy production would change from the spark-spread to the clean spread.
- Commodity prices (input fuel, output fuel) are stochastic.
- Must take into account (dynamic) strategies of other participants.
- Feedback between production policies and carbon prices.

Related Literature

- Real Options: Dixit and Pindyck (1994), Dockner et al. (2000).
- Analysis of Cap-and-Trade: Carmona et al. (2008,2009), Cetin and Verschuere (2008), ...
- Optimal Switching Problems (single-agent): Zervos (2003), Hamadène and Jeanblanc (2005), M.L. and Carmona (2008, 2009), Hu and Tang (2008).
- Optimal Stopping Games: Ohtsubo (1987,1991), Shmaya and Solan (2004), Ferenstein (2005,2007), Ramsey and Szajowski (2008).
- Stochastic Differential Games: Bensoussan, Friedman, Hamadène, Lepeltier,...

Model Setup

- We focus on the timing optionality within a real-options framework.
- Consider a duopoly – two producers (representing different sets of power plants).
- Each one sells electricity into a stochastic market at price P_t.
- Need emission permits to produce. Must buy CO_2 permits on the market at price X_t.
- Take a reduced-form price-impact model for (X_t) (do not explicitly model the remaining supply of permits).
- Simplify the strategy set: at each time epoch either produce, or stay offline, $\xi_i(t) \in \{0, 1\}$.
- Each producer's policy influences changes in X; $\Rightarrow$ the scheduling decisions of agents affect each other.
- Discrete-time model.

Basic Ingredients

- (P_t) is exogenously given and is a discrete-time process with Gaussian increments.
- At each instant t player i chooses emissions regime: $\xi_i(t)$.
- (X_t) is another 1-dim. process, drift is controlled by $\vec{\xi}(t)$; Gaussian increments correlated with those of (P_t).
- Changes in ξ_i are expensive (fixed switching costs $K_{i,\xi_i(t-),\xi_i(t)}$) and induce inertia and hysteresis.
- Net revenue is $\psi_i(p, x, \vec{\xi}) = (a_i P_t - b_i X_t - c_i)\xi_i(t)$.
- Fixed horizon T: expiration date of the permits.
- Each admissible control pair $\vec{\xi}$ induces a probability law $\mathbb{P}^{\vec{\xi}}$ of (X_t) through the price impact mechanism. Work with the physical measures.
- Each producer optimizes

$$V^i(0, p, x, \vec{\zeta}) = \mathbb{E}^{\vec{\xi}}\left[\sum_{t=0}^{T-1}\left(\xi_i(t)(a_i P_t - b_i X_t^{(\xi)} - c_i) - K_{i,\xi_i(t-),\xi_i(t)}\right)\right].$$

Dynamic Decision-Making

- At each step, each producer decides whether to produce or not.
- The chosen action results in immediate date t-payoff, as well as different continuation values on $[t+1, T]$.
- Leads to a repeated 2×2 stochastic game.
- Bellman's Principle is replaced by a game Nash Equilibrium (NE).
- Pure Nash equilibria might not exist.
- Existence: Need mixed equilibria.
- Might also have multiple Nash equilibria.
- Uniqueness: equilibrium refinement.

Dynamic Decision-Making

- At each step, each producer decides whether to produce or not.
- The chosen action results in immediate date t-payoff, as well as different continuation values on $[t+1, T]$.
- Leads to a repeated 2×2 stochastic game.
- Bellman's Principle is replaced by a game Nash Equilibrium (NE).
- Pure Nash equilibria might not exist.
- Existence: Need mixed equilibria.
- Might also have multiple Nash equilibria.
- Uniqueness: equilibrium refinement.

Dynamic Decision-Making

- At each step, each producer decides whether to produce or not.
- The chosen action results in immediate date t-payoff, as well as different continuation values on $[t+1, T]$.
- Leads to a repeated 2×2 stochastic game.
- Bellman's Principle is replaced by a game Nash Equilibrium (NE).
- Pure Nash equilibria might not exist.
- Existence: Need mixed equilibria.
- Might also have multiple Nash equilibria.
- Uniqueness: equilibrium refinement.

Multiple Local Equilibria

- Payoff functions: $\psi_1 = P_t - 2X_t - 10$ (coal plant);
  $\psi_2 = 2P_t - X_t - 80$ (clean gas plant).
- The value functions are decreasing and convex in x around $X_t = 10$.
- Today $P_t = 50, X_t = 15$, both producers are offline but in-the-money. Tomorrow $\mathbb{E}[P_{t+1}] = 50$.
- Net profit is $\Delta V_i = \mathbb{E}^{\vec{\xi}}[V_i(P_{t+1}, X_{t+1})] - V_i(p,x) + \psi_i(p,x)$.
- Expectations for tomorrow

Strategy $\vec{\xi}$	$\mathbb{E}^{\vec{\xi}}[X_{t+1}]$	Net $\Delta \vec{V}$
$(0,0)$	15	$(0,0)$
$(0,1)$	19	$(-3,3)$
$(1,0)$	23	$(2,-5)$
$(1,1)$	27	$(-5,-6)$

Classification of 2×2 Games

Three equivalence classes:

- A single dominating pure equilibrium (unanimity).
- A competitive game (essentially zero-sum) which admits a unique mixed Nash eqm.
- A (anti-) coordination game which admits two pure eqm's, a mixed one and a continuum of correlated eqm's – *"battle-of-the-sexes"* as above.
- Profitable for each one to emit separately; not profitable to emit together.
- Which producer will yield??

Correlated Equilibria

- Nash equilibrium: given the eqm strategy of the other player, maximizes your expected payoff.
- Overall payoff distribution is a product measure on the payoff space.
- A correlated equilibrium (γ^{jk}) is a general probability distribution on the payoff space. Known to all.
- Achieved by introducing a third (fictitious) agent, (regulator).
- The regulator sends a private signal $\mu_i(\gamma) \in \{0, 1\}$ to player i.
- Given the signal (and implied strategy of the second player), no incentive to deviate.

Formal 2×2 Game

- Payoffs $H = \begin{pmatrix} (\alpha^{00}, \beta^{00}) & (\alpha^{01}, \beta^{01}) \\ (\alpha^{10}, \beta^{10}) & (\alpha^{11}, \beta^{11}) \end{pmatrix}$.
- A policy is $(\vec{\pi}, \vec{\rho})$ whence π_i (resp. ρ_j) is the prob. that player 1 (player 2) chooses action i.
- Value of a policy to players is $Val(H; \vec{\pi}, \vec{\rho}) := \begin{pmatrix} \sum_{i,j} \pi_i \rho_j \alpha^{ij} \\ \sum_{i,j} \pi_i \rho_j \beta^{ij} \end{pmatrix}$.
- $\gamma = (\gamma^{ij})$ is a CE if

$$\begin{cases} \gamma^{00}\alpha^{00} + \gamma^{01}\alpha^{01} \geq \gamma^{00}\alpha^{10} + \gamma^{01}\alpha^{11}, & \gamma^{11}\alpha^{11} + \gamma^{10}\alpha^{10} \geq \gamma^{11}\alpha^{01} + \gamma^{10}\alpha^{00} \\ \gamma^{00}\beta^{00} + \gamma^{10}\beta^{10} \geq \gamma^{00}\beta^{01} + \gamma^{10}\beta^{11}, & \gamma^{11}\beta^{11} + \gamma^{01}\beta^{01} \geq \gamma^{11}\beta^{10} + \gamma^{01}\beta^{00}. \end{cases}$$

- Leads to game values $Val_\gamma(H) := \begin{pmatrix} \sum_{i,j} \gamma^{ij}\alpha^{ij} \\ \sum_{i,j} \gamma^{ij}\beta^{ij} \end{pmatrix}$.

More on Correlated Equilibria

- The set of correlated equilibria includes the convex hull of all Nash equilibria.
- The correlation mechanism must be known in advance – in a multistage game $\gamma(t)$ can be Markovian in state variables.
- Common choices:
  - A utilitarian mechanism, maximizing sum of game value for the agents.
  - An egalitarian mechanism, maximizing the minimum game value of the agents;
  - A fixed preferential mechanism, maximizing the game value for player 1 (resp. player 2);
- Given the signal, all actions are pure.

Stopping Games

- A stopping game: each agent chooses a stopping time τ_i, $i = 1, 2$. Stopping corresponds to action '1'.
- Payoff structure $(\mathcal{Z})$; agent i receives ($\tau \equiv \tau_1 \wedge \tau_2$)

$$\left(\sum_{t=0}^{\tau-1} Z_i^{00}(t)\right) + Z_i^{10}(\tau)1_{\{\tau_i < \tau_j\}} + Z_i^{01}(\tau)1_{\{\tau_i > \tau_j\}} + Z_i^{11}(\tau)1_{\{\tau_i = \tau_j\}}.$$

- Starting with known values at T move back in time; each period yields a 2-by-2 game with payoffs corresponding to conditional expectation of next-period value.
- Let $Val_\gamma(\mathcal{Z}_t)$ be an equilibrium of a 2-by-2 one-period game with payoffs

$$\mathcal{Z}_t = \begin{pmatrix} (\tilde{Z}_1(t), \tilde{Z}_2(t)) & (Z_1^{01}(t), Z_2^{01}(t)) \\ (Z_1^{10}(t), Z_2^{10}(t)) & (Z_1^{11}(t), Z_2^{11}(t)) \end{pmatrix}.$$

- Stopping game values solve $(V_1(t), V_2(t)) = Val_\gamma(\mathcal{Z}_t)$, with $(\tilde{Z}_1(t), \tilde{Z}_2(t)) \equiv (\mathbb{E}[V_1(t+1)|\mathcal{F}_t] + Z_1^{00}(t), \mathbb{E}[V_2(t+1)|\mathcal{F}_t] + Z_2^{00}(t))$.

Stopping Games

- A stopping game: each agent chooses a stopping time τ_i, $i = 1, 2$. Stopping corresponds to action '1'.
- Payoff structure $(\mathcal{Z})$; agent i receives $(\tau \equiv \tau_1 \wedge \tau_2)$

$$\Big(\sum_{t=0}^{\tau-1} Z_i^{00}(t)\Big) + Z_i^{10}(\tau)1_{\{\tau_i<\tau_j\}} + Z_i^{01}(\tau)1_{\{\tau_i>\tau_j\}} + Z_i^{11}(\tau)1_{\{\tau_i=\tau_j\}}.$$

- Starting with known values at T move back in time; each period yields a 2-by-2 game with payoffs corresponding to conditional expectation of next-period value.
- Let $Val_\gamma(\mathcal{Z}_t)$ be an equilibrium of a 2-by-2 one-period game with payoffs

$$\mathcal{Z}_t = \begin{pmatrix} (\tilde{Z}_1(t), \tilde{Z}_2(t)) & (Z_1^{01}(t), Z_2^{01}(t)) \\ (Z_1^{10}(t), Z_2^{10}(t)) & (Z_1^{11}(t), Z_2^{11}(t)) \end{pmatrix}.$$

- Stopping game values solve $(V_1(t), V_2(t)) = Val_\gamma(\mathcal{Z}_t)$, with $(\tilde{Z}_1(t), \tilde{Z}_2(t)) \equiv (\mathbb{E}[V_1(t+1)|\mathcal{F}_t] + Z_1^{00}(t), \mathbb{E}[V_2(t+1)|\mathcal{F}_t] + Z_2^{00}(t))$.

Nash Equilibria in Stopping Games

- To show existence of a pure Nash equilibrium need restrictive assumptions (e.g. Dynkin zero-sum games, non-zero-sum monotone games where $Z_i^{01} \le Z_i^{11} \le Z_i^{10}$).
- In general, must allow randomized stopping times.
- This is an $(\mathcal{F}_t)$-adapted stochastic process $p = (p_t)$ with $0 \le p_t \le 1$ a.s.
- $\tau(p) \triangleq \inf\{t : \eta_t \le p_t\}$ where $\eta_t \sim Unif(0,1)$ i.i.d.. p_t is the probability of stopping at date t, conditional on not stopping so far.
- $\tau(p)$ is not $(\mathcal{F}_t)$-adapted. Enlarge the filtration: $\tau(p)$ is a $(\mathcal{F}_t \vee \sigma(\eta_t))$-stopping time.
- Shmaya & Solan (2004): any discrete-time stopping game admits a mixed NE.
- Ferenstein (2005) gave a construction using backward recursion. Solution relies on recursive Nash equilibria and conditional expectations.

CE in Stopping Games

- A correlation law $(\gamma^{jk}(t))$ is a function of (t, P_t, X_t) and fixed/known in advance. Gives a CE for any payoff structure $\mathcal{Z}_t$.
- At each state t, agent i receives a private signal $\mu_i(t; \gamma)$.
- Resulting randomized stopping time is $\tau_i(\gamma)$. τ_i, τ_j are dependent!
- At each stage $\gamma^{jk}(t)$ induces a CE – no incentive to deviate given $\mu_i(t; \gamma)$.
- Admissible overall strategies are $\mathcal{G}^i$-stopping times τ_i, with $\mathcal{G}^i_t = \sigma(P_s, X_s, \mu_i(s), 0 \le s \le t)$.
- Game is non-cooperative; no possibility of threats, etc. Even if deviate continue to receive future messages and no changes are made.

Switching Game I

- We have a switching game. This is a sequential stopping game: can repeatedly "stop" to alter production regimes in response to changing electricity prices, permit prices or other agent's actions.
- Player i: value function $V_i(t, P_t, X_t, \vec{\xi}_t)$.

$$(V_1(t,\vec{\zeta}), V_2(t,\vec{\zeta})) =$$

$$Val_\gamma \begin{pmatrix} (\mathbb{E}^{\vec{\zeta}}[V_1(t+1,\vec{\zeta})|\mathcal{F}_t] + Z_1(t), \mathbb{E}^{\vec{\zeta}}[V_2(t+1,\vec{\zeta})|\mathcal{F}_t] + Z_2(t)) & (V_1(t,\zeta_1,\bar{\zeta_2}), V_2(t,\zeta_1,\bar{\zeta_2}) - K_{2,\zeta_2}) \\ (V_1(t,\bar{\zeta_1},\zeta_2) - K_{1,\zeta_1}, V_2(t,\bar{\zeta_1},\zeta_2)) & (V_1(t,\bar{\zeta_1},\bar{\zeta_2}) - K_{1,\zeta_1}, V_2(t,\bar{\zeta_1},\bar{\zeta_2}) - K_{2,\zeta_2}) \end{pmatrix}$$

where $Z_i(t) \triangleq (a_i P_t - b_i X_t - c_i)\zeta_i$.

- Overall structure:
  - Observe current state $(P_t, X_t, \vec{\xi}_{t-1})$;
  - Regulator carries out randomization;
  - Receive private signals $\mu_i(t;\gamma)$;
  - Choose private actions;
  - Joint action $\vec{\xi}_t$ is revealed, update state variables for next period;

Switching Game II

- Sketch of proof: Restrict strategy sets so that agents can only use up to (n, m) switches.
- Translates into an iterative stopping game with payoffs corresponding to $(n-1, m)$, $(n, m-1)$ or $(n-1, m-1)$ cases.
- Fixing the strategy of one player; the other player solves a switching problem with respect to the enlarged filtration $\mathcal{G}^i$.
- By definition of γ this gives a CE in the switching game.
- Take $n, m \to \infty$ to obtain a coupled pair of value functions as above.

Digression: Single Player Case

- Fix the strategy of one producer and consider the optimization of the other one.
- This becomes an optimal switching problem as studied in Carmona-M.L. (2008).
- The price impact leads to significant hysteresis effect.
- If the price impact is severe enough, will always stay offline (or at least with very high probability) – "blockading".
- From player's 1 perspective, the actions of player 2 are randomized: continuation values are unknown, optimal stopping in "random environment".
- Otherwise, standard optimal stopping problem in the enlarged filtration $(\mathcal{G}^i_t)$ that incorporates CE.

Numerical Solution

- To solve for the game values numerically need to
  - Be able to compute equilibria in 2×2 games;
  - Compute conditional expectations.
- Have explicit formulas for CE of 2×2 games (answer depends on CE choice).
- Need approximation; recall that (P, X) have continuous space.
- Need to work with four different prob. measures $\mathbb{P}^{\vec{\zeta}}$ due to the price impact.

Markov Chain Approximation

- Method I: discretize the state space of (P, X).
- Choose finite-state Markov $(\tilde{P}, \tilde{X})$ whose 1-step transitions are consistent with those of (P, X).
- Take a rectangular grid; allow $(\tilde{P}, \tilde{X})$ to have transitions only to neighboring grid points;
- Conditional expectations reduce to weighted sums;
- See the book by Kushner and Dupuis (2001).
- Generic proof of convergence of this approach for finite-action non-zero-sum stochastic game with 2 players was done in Kushner (2007).

Least Squares Monte Carlo

- To compute the conditional expectations, another robust algorithm is to use Monte Carlo simulation.
- Simulate paths of (P, X) for each of the four possible emission regimes $\vec{\zeta}$.
- Continuation values are approximated through a cross-sectional regression.
- If the optimal decision is to switch to another regime, then use the approximate continuation value; else recursively update the future path-value.
- Extends the Longstaff-Schwartz method for American option pricing (a single optimal stopping problem).
- A single-agent switching problem was solved in Carmona-M.L. (2008).
- Straightforward extension to randomized stopping ... and to 2-player game.

More on Simulation

- Main challenge: to account for price-impact;
- Values of X_t depend on chosen (in the past) $\vec{\xi}(t)$, but the algorithm is backward in time;
- Solution: re-simulate forward X_t paths based on known forward game values; update continuation values as realized payoffs in a simulated sub-game on $[t, T]$.
- Everything must be done for each possible future regime $\vec{\zeta}$.

Example

- $P_{t+1} = P_t \cdot \exp(2(50 - \log P_t) \ + 0.4\epsilon_t^P);$
- $X_{t+1} = X_t \cdot \exp(3(\log(12 + 8\xi_1(t) + 4\xi_2(t) - \log X_t) + 0.25\epsilon_t^X)$ with $\mathbb{E}[\epsilon^P \epsilon^X] = 0.6;$
- Revenues: $Z_1(t) = P_t - 2X_t - 10;$
  $Z_2(t) = 2P_t - X_t - 80;$
- $T = 1$, 26 periods ($\Delta t = 1/26$); $K \equiv 0.2$.
- Using the simulation solver:
-
Correlation Law	$V_1(0, P_0, X_0)$	$V_2(0, P_0, X_0)$
Utilitarian	5.30	4.14
Egalitarian	5.33	4.20
Preferential 1	5.39	4.11
Preferential 2	5.02	4.24

Local Equilibria

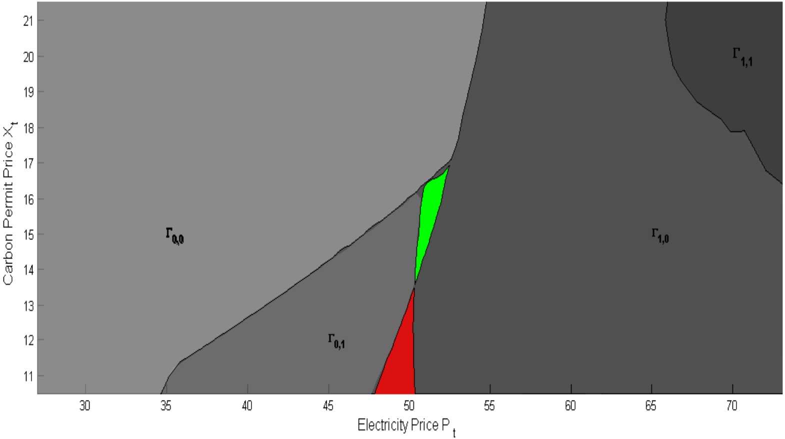

Optimal game strategy ξ^* as a function of (P_t, X_t) for $t = 0.25$. Here $\vec{\zeta} = (0, 0)$. The green region denotes the anti-coordination CE and the red region denotes the competitive mixed NE.

A Realized Equilibrium Path

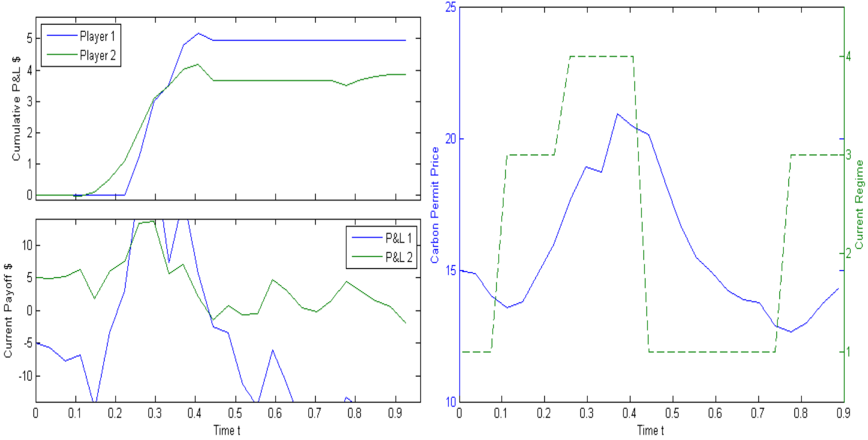

Sample path of the controlled X_t, including the corresponding strategy $\xi^* \in \{00, 01, 10, 11\}$. The top left panel shows the cumulative P&L of each player; the bottom left panel shows the raw P&L for each time period. Finally, the right panel shows the evolution of the controlled X_t, as well as the implemented strategy (ξ_t^1, ξ_t^2). Note as ξ_t increases, emissions rise and X_t tends to increase.

Conclusion

- Stochastic games naturally occur in studying oligopolies.
- The emission market would be a new important class of such problems.
- Investigate the simplest possible scenario where the game is non-trivial: a new model of an optimal switching game.
- Already the problems of equilibrium-refinement and computational tractability arise.

Open Problems

- The permit price is endogenous to the duopoly problem – current price should be (a function of) conditional expectation of total future emission of CO_2.
- Are there any no-arbitrage restrictions in this market (depends on what financial strategies participants may use)?
- Extend to a general equilibrium setup.
- Also, producers will be allowed to bank and trade their permits. How to incorporate initial permit allocations?

Continuous Time Formulation

- Would like a continuous-time counterpart of the discrete-time model.
- There are no existing results on general sequential stochastic games.
- Here we have a natural structure – switching game as a sequence of stopping games.
- Need to define/understand continuous-time CE?
- Randomized stopping times were considered in zero-sum context by Touzi and Vieille (2002).
- In our diffusion setup, should be related to solution of (obliquely reflected) high-dimensional BSDE.
- Such representations have been obtained in related stochastic differential games or zero-sum Dynkin games.
- Reflection rule is complicated – uses 8 different “continuation” processes!

References I

California Air Resources Board
Climate Change Proposed Scoping Plan.
http://www.arb.ca.gov, October 2008.

E. J. Dockner, S. Jorgensen, N.V. Long, and G. Sorger,
Differential games in economics and management science.
Cambridge University Press, Cambridge, 2000.

R. Carmona, M. Fehr, J. Hinz and A. Porchet.
Market Design for Emission Trading Schemes,
SIAM Review, forthcoming.

S. Hamadéne and J. Zhang.
The continuous time nonzero-sum Dynkin game problem and application in game options, 2009.
prepring, available at http://www-rcf.usc.edu/ jianfenz/Papers/HZ2.pdf.

H. J. Kushner.
Numerical approximations for nonzero-sum stochastic differential games.
SIAM J. Control Optim., 46(6):1942–1971, 2007.

Y. Ohtsubo.
A nonzero-sum extension of Dynkin's stopping problem.
Math. Oper. Res., 12(2):277–296, 1987.

D. M. Ramsey and K. Szajowski.
Selection of a correlated equilibrium in Markov stopping games.
European J. Oper. Res., 184(1):185–206, 2008.

References II

R. Carmona, and M. Ludkovski.
Pricing asset scheduling flexibility using optimal switching.
Appl. Math. Finance, 15(4):405–447, 2008.

M. Ludkovski
Stochastic Switching Games and Duopolistic Competition in Emissions Markets.
In preparation, available on request.