House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

A sequential design for gas storage optimization using kriging metamodels

Séminaire Chaire FDD et IdR FiME

Stephan Prell | Chair for Energy Trading and Finance | University of Duisburg-Essen

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Agenda

Introduction

Review of Markov decision processes

The storage optimization problem and the RMC approach

Kriging

Sequential design augmenting

Illustration

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Motivation

- We are considering an optimal control problem with a (small) finite action space which is a common problem in energy markets.
- An area where this kind of problems arise frequently, is the valuation of operational flexibility of exotic energy derivatives (e.g. power plants, swing options, hydro pumped storage, batteries)
- Problems with larger action spaces than $\mathfrak{L} = \{\mathit{continue}, \mathit{stop}\}$ are sometimes referred to as optimal switching problems
- These problems in energy markets are sometimes characterized by idiosyncratic price processes and high dimensional state spaces.

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Motivation

- We try to develop an adaptive design to optimize a gas storage in the backward dynamic programming framework using Regression Monte Carlo
- LSMC applied to American options is studied in Longstaff and Schwartz (2001), Tsitsiklis and Van Roy (2001) and many more. . .
- Boogert and De Jong (2008) and Carmona and Ludkovski (2010) used LSMC for valuation of gas storages
- Kohler (2010) reviews more general RMC frameworks for pricing American options
- Recently Gramacy and Ludkovski (2015) and Ludkovski (2016) studied sequential sampling for RMC for optimal stopping problems and Hu and Ludkovski (2016) for ranking response surfaces
- We aim to develop an algorithm for optimal switching problems, where a priori little is known about the exact decision boundary

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Controlled Markov process

- We consider random dynamics on a finite time horizon $0, \ldots, T$ whose state x evolves in $\mathcal{X}$ and is controlled by actions j from a finite action set $\mathfrak{L}$
- $\mathcal{D}_t \subset \mathcal{X} \times \mathfrak{L}$ denotes the set of possible state-action combinations at time t
- We denote $\pi_t : \mathcal{X} \to \mathfrak{L}$ a decision rule $((x, \pi_t) \in \mathcal{D}_t)$ and a sequence of decision rules $\pi = (\pi_t)_{t=0}^{T-1}$ a strategy
- For each $j \in \mathfrak{L}$, $\mathcal{K}_t^j(x, dx')$ is a stochastic transition kernel from $\mathcal{D}_t$ to $\mathcal{X}$, i.e. for each initial point $x_0 \in \mathcal{X}$ and each policy π there exist a probability measure $\mathbb{P}^{x_0,\pi}$ and a stochastic process $(X_t)_{t=0}^T$ satisfying the initial condition $\mathbb{P}^{x_0,\pi}(X_0 = x_0) = 1$ such that

$$\mathbb{P}^{x_0,\pi}(X_{t+1} \in \mathcal{B} | X_0, \ldots, X_t) = \mathcal{K}^{\pi_t(X_t)}(X_t, \mathcal{B})$$

holds for each $\mathcal{B} \subset \mathcal{X}$ for all t.

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Controlled Markov process

- For an action $j \in \mathfrak{L}$, we use $\mathcal{K}_t^j$ as the one-step transition operator , who acts on functions v by

$$(\mathcal{K}_t^j v)(x) = \int_{\mathcal{X}} v(x')\mathcal{K}_t^j(x, dx')$$

  whenever the integrals are well defined
- $\psi_j(t,x) : \mathcal{D}_t \to \mathbb{R}$ gives the (discounted) one-stage reward
- $\psi(T,x) : \mathcal{X} \to \mathbb{R}$ is the scrap function for the reward at T
- Having an initial point x_0 and a stochastic process, our goal is to maximize the expected total reward, i.e. to find the optimal strategy $\pi^* = (\pi^*)_{t=0}^{T-1}$ such that

$$\pi^* = \arg\max_{\pi} \mathbb{E}^{x_0,\pi}\left(\sum_{t=0}^{T-1} \psi_{\pi(X_t)}(t, X_t) + \psi(T, X_t)\right)$$

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Controlled Markov process

- To obtain the optimal strategy, we introduce for $t, \ldots, T-1$ the maximal reward operator

$$\Gamma_t v(x) = \sup_{j \in \mathcal{D}_t} (\psi_j(t,x) + \mathcal{K}_t^j v(x))$$

  which acts on each measurable function $v : \mathcal{X} \to \mathbb{R}$
- We consider the Bellman recursion

$$v_T^* = \psi(T, \cdot) \qquad\qquad v_t^* = \Gamma_t v_{t+1}^*$$

- There exist a recursive solution to the Bellman recursion, which gives the value functions and determines an optimal strategy π^*

$$\pi_t^*(x) = \arg\max_{j \in \mathcal{D}_t} (\psi_j(t,x) + \mathcal{K}_t^j v_{t+1}^*(x))$$

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Control problem of a storage

- Given a finite time horizon $\{0, 1, \ldots, T\}$, we consider a controlled Markovian process $(X_t^j)_t^T := (C_t, \mathcal{G}_t, i)$, where
  - $c_{\min} \leq C_t \leq c_{\max} \ \forall t$ is the controlled inventory process
  - $\mathcal{G}_t = \{\mathcal{G}_t^p\}_{p=1}^d$) are d uncontrolled stochastic processes realized on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, which drive the price process $G_t(\mathcal{G}_t) : \mathbb{R}^d \to \mathbb{R}$
  - $i \in \mathfrak{L} = \{inject, store, withdraw\}$ is the operating regime
- Furthermore we have a cost component $K_{i,j}$ for switching the system from i to j
- A constant discount rate r
- The maximal rate, at which we can change the inventory upwards $\bar{a}_{\text{in}}(C_t)$ and downwards $\bar{a}_{\text{out}}(C_t)$ s.t. the constrains
- Buying and selling at the market is represented by $\psi_{j_t}(G_t(\mathcal{G}_t), C_t)$, which is linear in $X_t^{j_t}$ and also includes other storage costs for applying action j_t at time t.

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Control problem of a storage

- Having a linear reward function we know $C'_t(C_t, j_t)$, the next state of C_t if we a apply j_t, due to the 'bang-bang' property
- Provided with a scrap function, we find the solution of the Bellman recursion by the value functions

$$V(t, \mathcal{G}_t, C_t, i) = \max_{j_t \in \mathcal{D}_t} [-K_{i,j_t} + \psi_{j_t}(G_t(\mathcal{G}_t), C_t) + e^{-r}\mathbb{E}\left(V(t+1, \mathcal{G}_{t+1}, C'(C_t, j_t), j_t)|\mathcal{F}_t\right)],$$

where we suppressed the time dependency in ψ and spread $X_t^{j_t} = C_t \times \mathcal{G}_t$.

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Control problem of a storage

- We define the continuation function as

$$\mathfrak{C}(t, \mathcal{G}_t, C'(C_t, j_t), j_t) := \mathbb{E}\left(V(t+1, \mathcal{G}_{t+1}, C'(C_t, j_t), j_t) | \mathcal{F}_t\right)$$

for each $(t, \mathcal{G}_t, C'(C_t, j_t), j_t)$ with $t \leq T-1$, and the (backwards) stochastic dynamic program continuation function

$$\begin{aligned}\mathfrak{C}(t, \mathcal{G}_t, C'(C_t, j_t), j_t) :=& \mathbb{E}\left(\max_{j_{t+1} \in \mathcal{D}_{t+1}} \left[-K_{j_t, j_{t+1}} + \psi_{j_{t+1}}(G_{t+1}(\mathcal{G}_{t+1}), C'(C_t, j_t))\right.\right. \\ &\left.\left. + e^{-r}\mathfrak{C}(t+1, \mathcal{G}_{t+1}, C'(C_{t+1}, j_{t+1}), j_{t+1})\right] \middle| \mathcal{F}_t\right)\end{aligned}$$

- We use regression Monte-Carlo to approximate the continuation function

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Regress-now / regress later monte-carlo

- In the classical regression Monte-Carlo a set of simulations and a scrap function are used to approximate the continuation function at time $t+1$ by a two-norm regression on a linear combination of $\mathfrak{K}$ basis functions $\phi_{t,\mathfrak{k}}, \quad \mathfrak{k} = 1, \ldots, \mathfrak{K}$ (regress-now Least Squares Monte-Carlo)
- Let $\Phi_{t,\mathfrak{k}}, \quad \mathfrak{k} = 1, \ldots, \mathfrak{K}$ be the $\mathfrak{k}$-th basis function for the value function approximation at time t. This so called regress-later Least Squares Monte-Carlo is based on the assumption, that for each basis function $\Phi_{t,\mathfrak{k}}$ the conditional expectation

$$\hat{\Phi}_{t,\mathfrak{k}}(\mathfrak{x}, C_{t+1}) = \mathbb{E}\left[\Phi_{t+1,\mathfrak{k}}(G_{t+1}(\mathcal{G}_{t+1}), C_{t+1}) | G_t(\mathcal{G}_t) = \mathfrak{x}\right]$$

  can be computed analytically (or precomputed numerically in practice), (see Nadarajah et al. (2017) for a comparison of regress-now / regress later)
- Choice of basis functions? Choice of $\mathfrak{K}$?

Value iteration / policy iteration

- Value iteration regression Monte Carlo for optimal stopping problems has been introduced by Tsitsiklis and Van Roy (2001)
- In terms of storage optimization, the continuation function estimate for the n-th sample path becomes

$$\mathfrak{C}^n(t,\mathcal{G}_t^n,C'(C_t^n,j_t^n),j_t^n) = \max_{j_{t+1}\in\mathcal{D}_{t+1}} \Big[-K_{j_t^n,j_{t+1}} + \psi_{j_{t+1}}(G_{t+1}(\mathcal{G}_{t+1}^n),C'(C_t^n,j_t^n)) + e^{-r}\hat{\mathfrak{C}}(t+1,\mathcal{G}_{t+1}^n,C'(C_{t+1}^n,j_{t+1}),j_{t+1})\Big]$$

- Policy iteration regression Monte Carlo for optimal stopping problems has been introduced by Longstaff and Schwartz (2001)
- In terms of storage optimization, the continuation function estimate for the n-th sample path becomes

$$\mathfrak{C}^n(t,\mathcal{G}_t^n,C'(C_t^n,j_t^n),j_t^n) = \max_{j_{t+1}\in\mathcal{D}_{t+1}} \Bigg[-K_{j_t^n,j_{t+1}} + \psi_{j_{t+1}}(G_{t+1}(\mathcal{G}_{t+1}^n),C'(C_t^n,j_t^n)) + \sum_{\bar{t}=t+2}^{T-1} e^{-(\bar{t}-(t-1))r} - K_{\hat{j}_{\bar{t}-1}^n,\hat{j}_{\bar{t}}^n} + \psi_{\hat{j}_{\bar{t}}^n}(G_{\bar{t}}(\mathcal{G}_{\bar{t}}^n),C'(C_{\bar{t}-1}^n,\hat{j}_{\bar{t}-1}^n)) + \psi(G_T(\mathcal{G}_T^n),C'(C_{T-1}^n,\hat{j}_{T-1}^n))\Bigg]$$

where

$$\hat{j}_t^n(j_{t-1}^n,\mathcal{G}_t^n,C_t^n) = \arg\max_{j_t\in\mathcal{D}_t} \Big[-K_{j_{t-1}^n,j_t} + \psi_{j_t}(G_{t+1}(\mathcal{G}_t^n),C'(C_{t-1}^n,j_{t-1}^n)) + e^{-r}\hat{\mathfrak{C}}(t,\mathcal{G}_t^n,C'(C_t^n,j_t),j_t)\Big]$$

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Optimal inventory control

- Contrary to the optimal stopping problem, the inventory process C_t depends on the control, so one cannot compute the approximation of $\mathfrak{C}(t, \mathcal{G}_t, C'(C_t, j_t), j_t)$ and then optimize over the controls.
- Solutions
  - Grid discretization of inventory levels (e.g. Boogert and De Jong (2008), Nadarajah et al. (2017))
  - Quasi-simulation of the inventory (e.g. Carmona and Ludkovski (2010))
  - Control Randomisation (e.g. Kharroubi et al. (2014))
  - Backward construction of inventory levels (e.g. Balata and Palczewski (2017))
  - Resimulation of the inventory levels

Regression techniques

- Different regression techniques to compute the approximation of $\mathfrak{C}(t, \mathcal{G}_t, C'(C_t, j_t), j_t)$
  - Ordinary least-squares regression on global polynomial basis functions (e.g. Boogert and De Jong (2008), Nadarajah et al. (2017), Cortazar et al. (2008), Carmona and Ludkovski (2010),Longstaff and Schwartz (2001), Balata and Palczewski (2017))
  - . . . supplemented with closed form option prices (e.g. Nadarajah et al. (2017))
  - Ordinary least-squares regression on weighted laguerre polynomials (Kiesel et al. (2010)) or Hermite, Hyperbolic and Chebyshev polynomials (e.g. Choudhury and King (2008))
  - Nonparametric regression / smoothing splines (see Kohler (2010))
  - Linear functions or polynomials with local hypercube support (e.g. Bouchard and Warin (2012))
  - Dynamic trees to partitioning the state space into the continuation and stopping regions (Gramacy and Ludkovski (2015))

Control mapping

- For the optimization of the storage we therefore have to rank the surfaces over the state space $\mathcal{X} \subseteq \mathbb{R}^{d+1}$. For any $x = (x_{g^1}, \ldots, x_c) \in \mathcal{X}$ we denote

$$Y_{t,i,j}(x) = -K_{i,j} + \psi_j(G(x_g), x_c) + e^{-r}\hat{\mathfrak{C}}(t, x_g, C'_t(x_c, j), j)$$

the (noisily sampled) estimated surfaces and

$$\mu_{t,i,j}(x) = -K_{i,j} + \psi_j(G(x_g), x_c) + e^{-r}V(t+1, x_g, C'(x_c, j), j)$$

the true surfaces

- The classifier

$$\mathcal{C}_{t,i}(x) = \arg\max_{j \in \mathfrak{L}} [\mu_{t,i,j}(x)]$$

determines the optimal switching strategy $\forall T \leq t \leq 0, \ \ i \in \mathfrak{L}$

Optimal switching

- Next, we introduce the classifier

$$\mathcal{A}_{t,i}(x) = \arg \max_{j \neq \mathcal{C}_{t,i}(x) \in \mathfrak{L}} [\mu_{t,i,j}(x)] \quad \forall\, T-1 \leq t \leq 0, \quad i \in \mathfrak{L},$$

  which gives the best alternative decision to the optimal switching decision.
- We use $\mathcal{A}_{t,i}(x)$ to define

$$\vartheta_{t,i}(x) = \mu_{t,i,\mathcal{C}_{t,i}(x)}(x) - \mu_{t,i,\mathcal{A}_{t,i}(x)}(x) \geq 0$$

  which is a measure for the 'triviality' of finding the optimal switching decision.
- Finding $\{x | \vartheta_{t,i}(x) = 0\}$ is a challenging task and together with $\mathcal{C}_{t,i}(x)$ the nested solution of the optimal control problem.

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Gaussian processes

- A metamodel is an approximation of the Input/Output function that is implied by the underlying simulation model.
- We want to use Gaussian processes to learn input-output mappings from the simulated data $\mathcal{E} = \{x^{1:n}, y^{1:n}\}$
- We use Gaussian processes to define distributions over functions $p(f)$ which can be used for (Bayesian) regression $p(f|\mathcal{E}) = \frac{p(f)p(\mathcal{E}|f)}{p(\mathcal{E})}$
- Gaussian processes can handle the case in which data is available in (multiple) different forms, as long as we can define an appropriate covariance function for each data type
- Let $\vec{f} = (f(x^1), f(x^2), \ldots, f(x^N))$ be an N-dimensional vector of function values evaluated at N points $x^i \in \mathcal{X}$

Definition

$p(f)$ is a Gaussian process if for any finite subset $\{x^1, \ldots, x^N\} \in \mathcal{X}$, the marginal distribution over that finite subset $p(\vec{f})$ has multivariate Gaussian distribution.

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Gaussian process regression: Kriging

- Kriging is used to build a global estimate of the entire response surface $\mathfrak{C}$
- We can use kriging for modeling of noise free experiments (interpolation) and stochastic simulators (smoothing)
- Assuming the model

$$y = f(x) + \epsilon$$

with $\epsilon \sim \mathcal{N}(0, \sigma^2)$, a $\mathcal{GP}$ prior and given $\mathcal{E}$ we can make predictions for new points x_*

$$p(y_*|x_*, \mathcal{E}) = \int p(y_*|x_*, f, \mathcal{E}) p(f, \mathcal{E}) df$$

- As the prior on f is a $\mathcal{GP}$ and the likelihood is Gaussian, the posterior on f is also a $\mathcal{GP}$
- Applying kriging we treat each $\mathfrak{C}$ as a realization of a Gaussian process specified by a trend function $f_j^*(x)$and a covariance structure $\mathcal{M}_j(x_u, x_v)$

Kriging

- For the covariance structure we choose a parametric kernel depending only on the increment $x_u - x_v = h = (h_1, \ldots, h_{d+1})$ and a multiplicative constant $\sigma^2 > 0$

$$\mathcal{M}(x_u, x_v) = \sigma^2 \prod_{l=1}^{d+1} g(h_l; \Theta_l)$$

where g is a 1-dimensional covariance kernel called Matern-5/2 kernel

$$g(h) = \left(1 + \frac{\sqrt{5}|h|}{\theta} + \frac{5h^2}{3\theta^2}\right) \exp\left(-\frac{\sqrt{5}|h|}{\theta}\right)$$

- Other choices include the Squared Exponential covariance function

$$g(h) = \exp\left(-\frac{h^2}{2\theta^2}\right)$$

- The only technical restriction on the covariance function is that it must be positive semi-definite

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Covariance functions

- The covariance function determines the smoothness of the functions

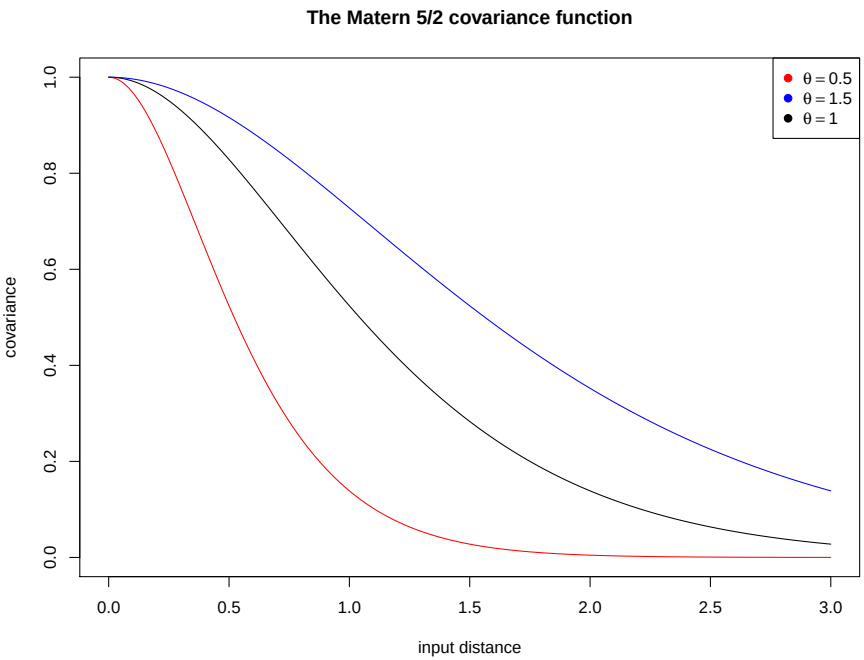

(a) Matern-5/2 covariance function

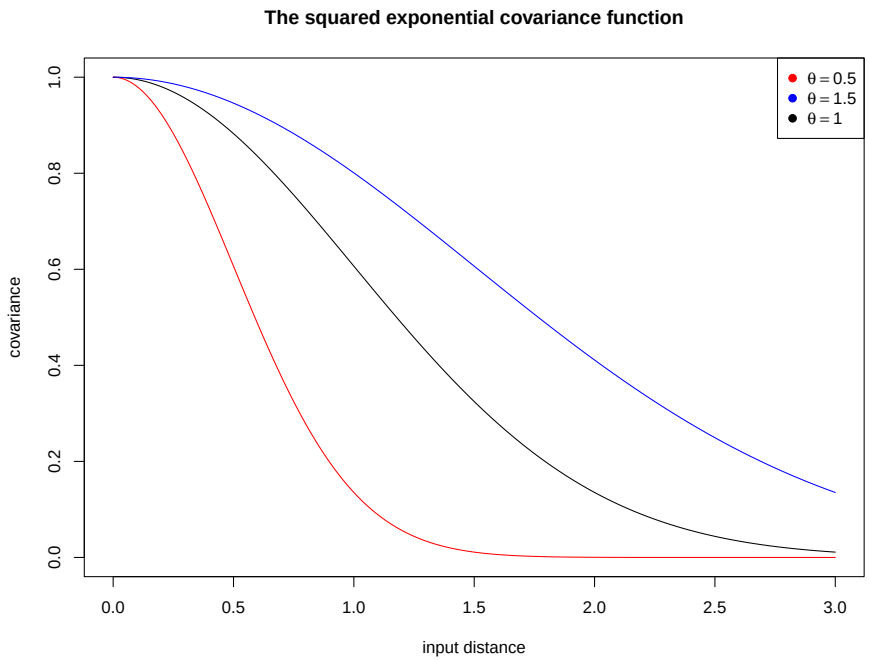

(b) Squared exponential covariance function

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Kriging

- Given the n noisy observations $\vec{y} = (y(x^1), \ldots, y(x^n))^T$ from the stochastic sampler at locations $\vec{x} = x^{1:n}$, the posterior of $\mathfrak{C}$ again forms a Gaussian process, specified by the mean m_j and the covariance s_j

$$m_j^{(n)}(x_u) = f_j^*(x_u) + \vec{k}_j^{(n)}(x_u)^T(\mathbf{K}_j + \Sigma_j^{(n)})^{-1}(\vec{y} - f_j^{*(n)}(\vec{x}))$$

$$s_j^{(n)}(x_u, x_v) = \mathcal{M}_j(x_u, x_v) - \vec{k}_j^{(n)}(x_u)^T(\mathbf{K}_j + \Sigma_j^{(n)})^{-1}\vec{k}_j^{(n)}(x_v)$$

where $\vec{k}_j^{(n)}(x_u) = (\mathcal{M}_j(x^1, x_u), \ldots, \mathcal{M}_j(x^n, x_u))^T$, $(\mathbf{K}_j)_{u,v} = \mathcal{M}_j(x_u, x_v)$, $1 \leq u, v \leq n$ and $\Sigma_j^{(n)} = \mathrm{diag}(\tau_j^2(x^1), \ldots, \tau_j^2(x^n))$ is the diagonal matrix of the noise variances

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Batching

- Inference on the kriging hyperparameters requires knowledge of the sampling noise
- We therefore sample M times at the same x and use the empirical estimates of τ^2,

$$\tilde{\tau}^2 = \frac{1}{M-1} \sum_{b=1}^{M} (y(x^b) - \bar{y}(x))^2$$

as proxy of the unknown τ^2, where $\bar{y}(x) = \frac{1}{M} \sum_{b=1}^{M} y(x^b)$ is our sample point $y(x)$

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Sequential design [Hu and Ludkovski (2016)]

- Instead of the static in sample regression, we want to develop a sequential design to learn the decision boundaries that determine the optimal switching strategy
- We assume sampling is not inexpensive and optimizing sampling efficiency is desirable
- Therefore we sample sequentially

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Loss function and local ranking complexity

- We want to assign the optimal strategy $\hat{\mathcal{C}}_{t,i}(x)$ to any $x \in \mathcal{X}$ that approximates the true $\mathcal{C}_{t,i}$ in terms of the loss function

$$\mathcal{L}_t(\hat{\mathcal{C}}_{t,i}, \mathcal{C}_{t,i}) = \mathbb{E}_t \left\{ \mu_{t,i,\mathcal{C}_{t,i}(x)}(x) - \mu_{t,i,\hat{\mathcal{C}}_{t,i}(x)}(x) \right\}$$

- Therefore sampling should be done in the regions, where identifying $\mathcal{C}_{t,i}(x)$ is difficult
- We measure the the local ranking complexity and the confidence in the estimated $\hat{\mathcal{C}}_{t,i}(x)$ by $\vartheta_{t,i}(x)$
- A low $\vartheta_{t,i}(x)$ indicates a difficult to determine switching decision, which requires high local accuracy in the estimation. Vice versa, a high $\vartheta_{t,i}(x)$ indicates a trivial decision.

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Sequential design

- We try to choose a collection of sample points, denoted by $\mathcal{Z} = (x, j)^{1:K}, x \in \mathcal{X}, \ j \in \mathfrak{L}$ such that we minimize the expected loss
- Having a design $\mathcal{Z}^k$ we iteratively add new design sites by optimization of an acquisition function that quantifies information gain
- Instead of optimizing the acquisition function over the state space directly, we propose an efficient approximation by selecting a set of candidate points $\mathcal{T}$ of size D via stochastic sampling
- Fixing t and i, we choose those $\bar{N} << D$ new points $x^{n+(1:\bar{N})} \in \mathcal{T}$, for which hold

$$m_{\hat{C}(x)}(x) - \zeta \sqrt{s_{\hat{C}(x)}(x,x)} < \max_{j \neq \hat{C}(x)} m_j(x) + \zeta \sqrt{s_j(x,x)}$$

where

$$\zeta = \min \left(\zeta \left| \sum_{l \in \mathcal{T}} \mathbb{1}_{\left\{ m_{\hat{C}(x^l)}(x^l) - \zeta \sqrt{s_{\hat{C}(x^l)}(x^l,x^l)} < \max_{j \neq \hat{C}(x^l)} m_j(x^l) + \zeta \sqrt{s_j(x^l,x^l)} \right\}} \geq \bar{N} \right)$$

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Computational burden reduction measures

- The computational bottleneck of the algorithm is the calculation of $\vec{k}_j^{(n)}(x)$ for the prediction of the posterior means
- Applying the batching schemes, we achieved the first reduction in the computational effort by reducing the length of $\vec{k}_j^{(n)}(x)$ to n, while having $n \times M$ samples.
- Without further measures, we still would have to calculate the covariance $n^2 \times M$ times per time step in the forward simulation.
- We can reduce the computational burden further by introducing a dense grid of points $\mathcal{S} \in \mathcal{X}$, at which we calculate $m_j(x)$ for all j and use $Y_{\cdot,i,j}(x)$ to find $\hat{\mathcal{C}}_{t,i}(x)$ for all $x \in \mathcal{S}$
- We apply a $p-$nearest neighbors algorithm in the forward simulation to reduce the amount of points at which we have to calculate $m_j(x)$ by the following principle: We calculate $m_j(x)$ only at those points, where the $p-$nearest neighbors of $\mathcal{S}$ are ambiguous about $\hat{\mathcal{C}}_{t,i}(x)$.
- Normally a choice of $p = 3^{(d+1)}$ gives admirable accuracy.

Updating Kriging

- Estimating the new model from scratch would be computationally inefficient. Fortunately kriging gives us efficient updating formulas

$$
\begin{aligned}
m_j^{(n+1)}(x_u) =& m_j^{(n)}(x_u) + \\
& \vec{k}_j^{(n+1)}(x_u)^T (\mathbf{K}_j + \Sigma_j^{(n+1)})^{-1} (y^{n+1} - m_j^{(n)}(x^{n+1})) \\
s_j^{(n+1)}(x_u, x_v) =& s_j^{(n)}(x_u, x_v) - (\vec{k}_j^{(n+1)}(x_u)^T (\mathbf{K}_j + \Sigma_j^{(n+1)})^{-1})^T \\
& (s_j^{(n)}(x^{n+1}, x^{n+1}) + \tau_j^2(x^{n+1})) \vec{k}_j^{(n+1)}(x_v)^T (\mathbf{K}_j + \Sigma_j^{(n+1)})^{-1}
\end{aligned}
$$

where $m_j^{(n+1)}(x_u)$ and $s_j^{2(n+1)}(x_u, x_v)$ is the new conditional posterior mean and covariance respectively.

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Illustration

As an illustrative example, we consider a gas storage with the following properties:

- $T = 1\text{year}, \Delta t = 1/100$
- $G_t = 17.1(\log 3 - \log G_t)dt + 1.33dW_t$
- $c_{\min} = 0Bcf, c_{\max} = 8Bcf$
- $a_{\text{in}}(C_t) = 0.09 \cdot 365\Delta t$, $a_{\text{out}}(C_t) = 0.2 \cdot 365\Delta t$
- $K_{i,j} = 0.25$ if $i \neq j$
- $$\psi_j(G_t, C_t) = \begin{cases} -a_{\text{in}}(C_t) \cdot G_t - 0.1\Delta t(C_t) & \text{if } j = \textit{inject} \\ -0.1\Delta t(C_t) & \text{if } j = \textit{store} \\ \bar{a}_{\text{out}}(C_t) \cdot G_t - 0.1\Delta t(C_t) & \text{if } j = \textit{withdrawal}. \end{cases}$$
- $r = 0.06$
- $V(T, g, c, i) = -2g \cdot \max[4 - c, 0]$
- D=2000
- $\bar{N} = 50$
- M=20
- $n = 15^2$

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Illustration for $t = 50$

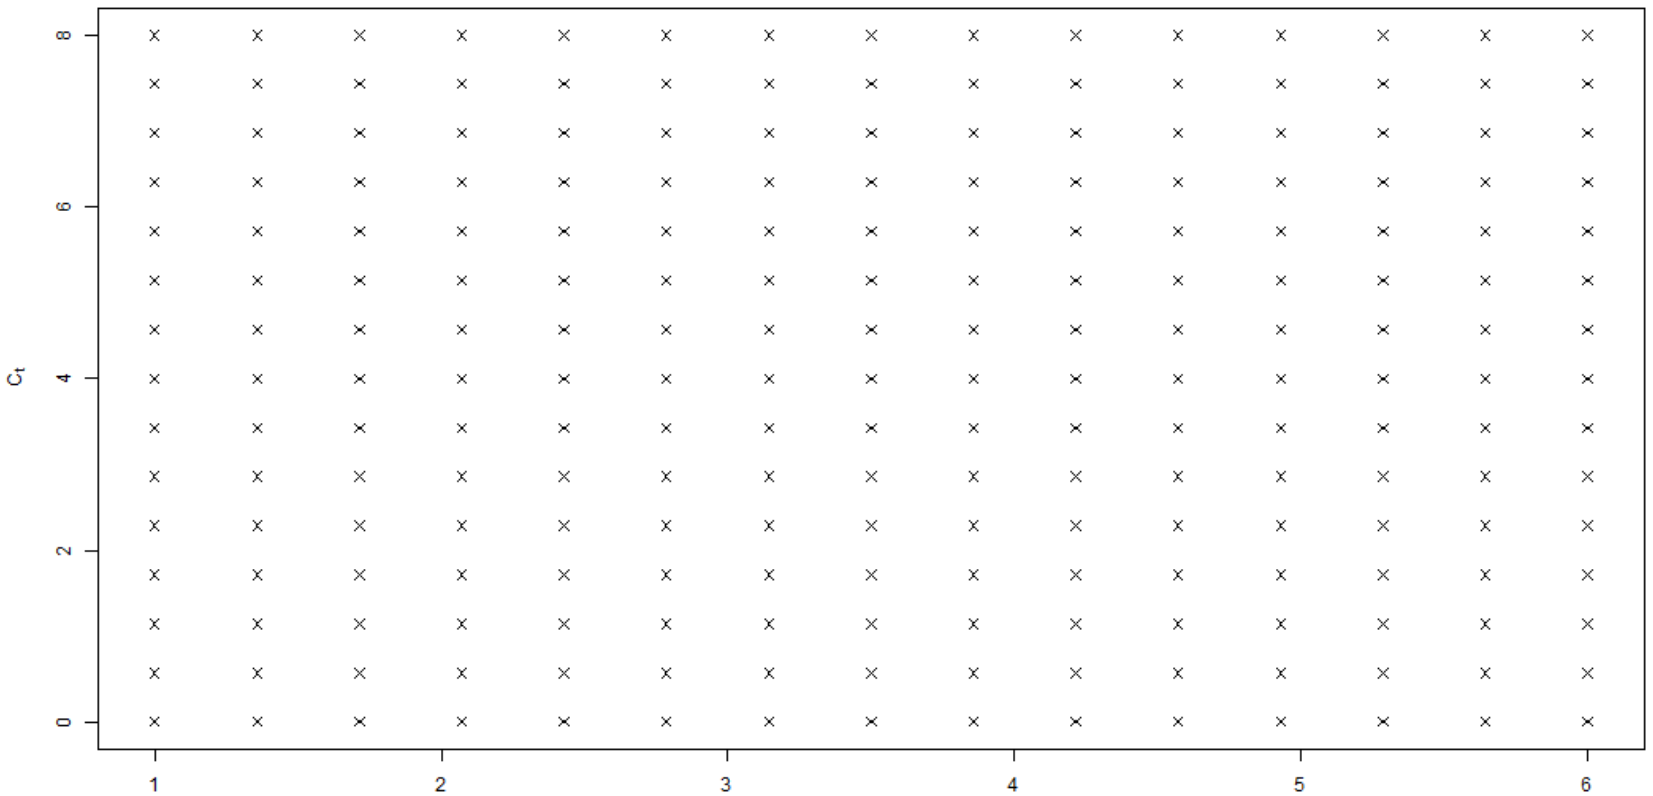

Figure: The grid $x^{(n)}$

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Illustration for $t = 50$

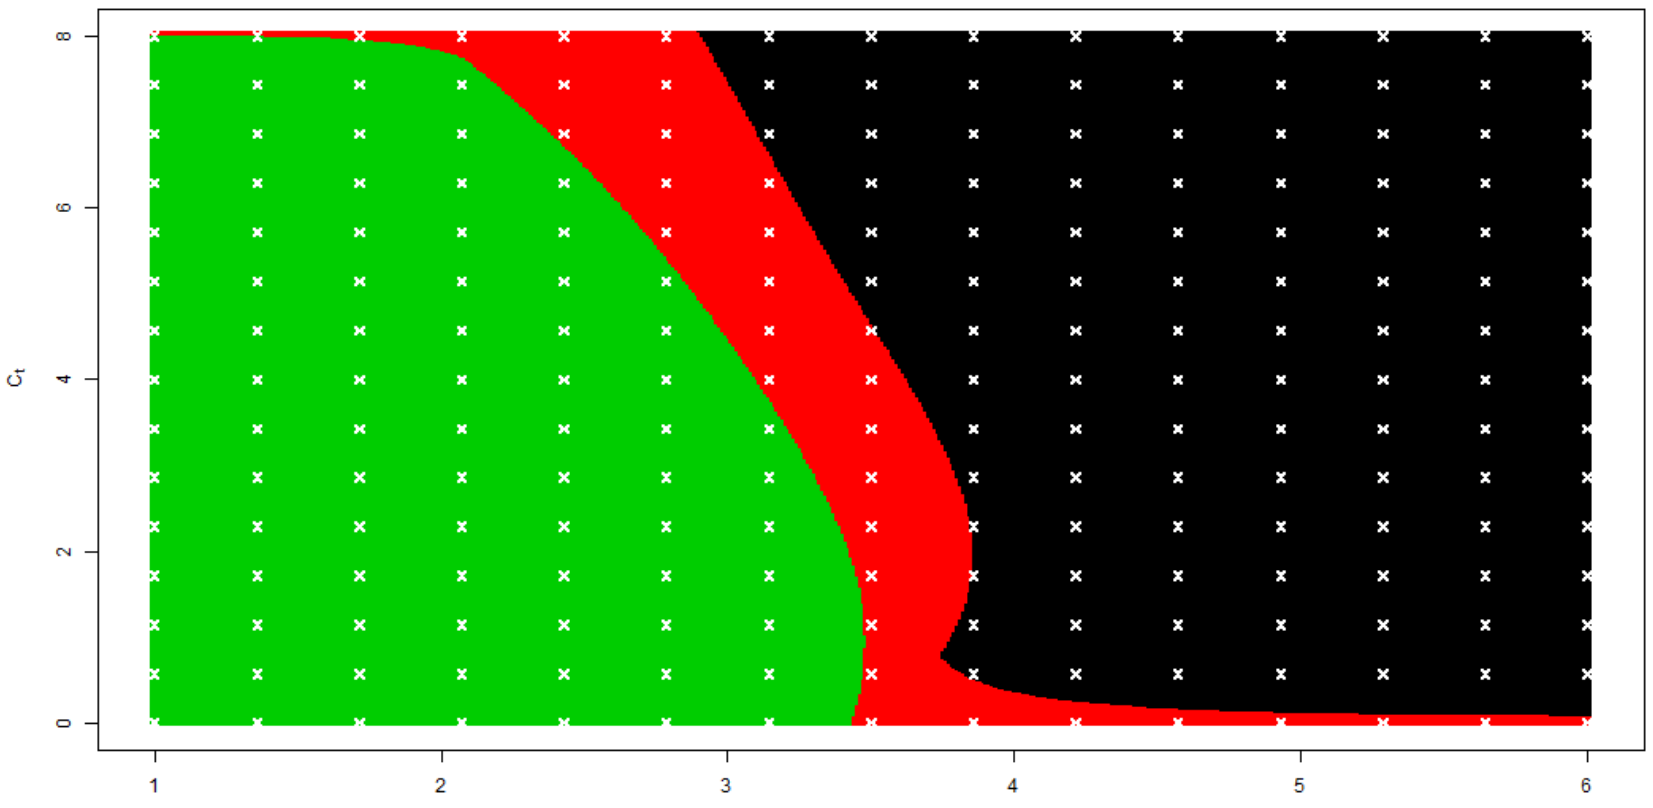

Figure: $\mathcal{C}_{50,store}(x)$,green:inject, red:store, black:withdraw

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Illustration for $t = 50$

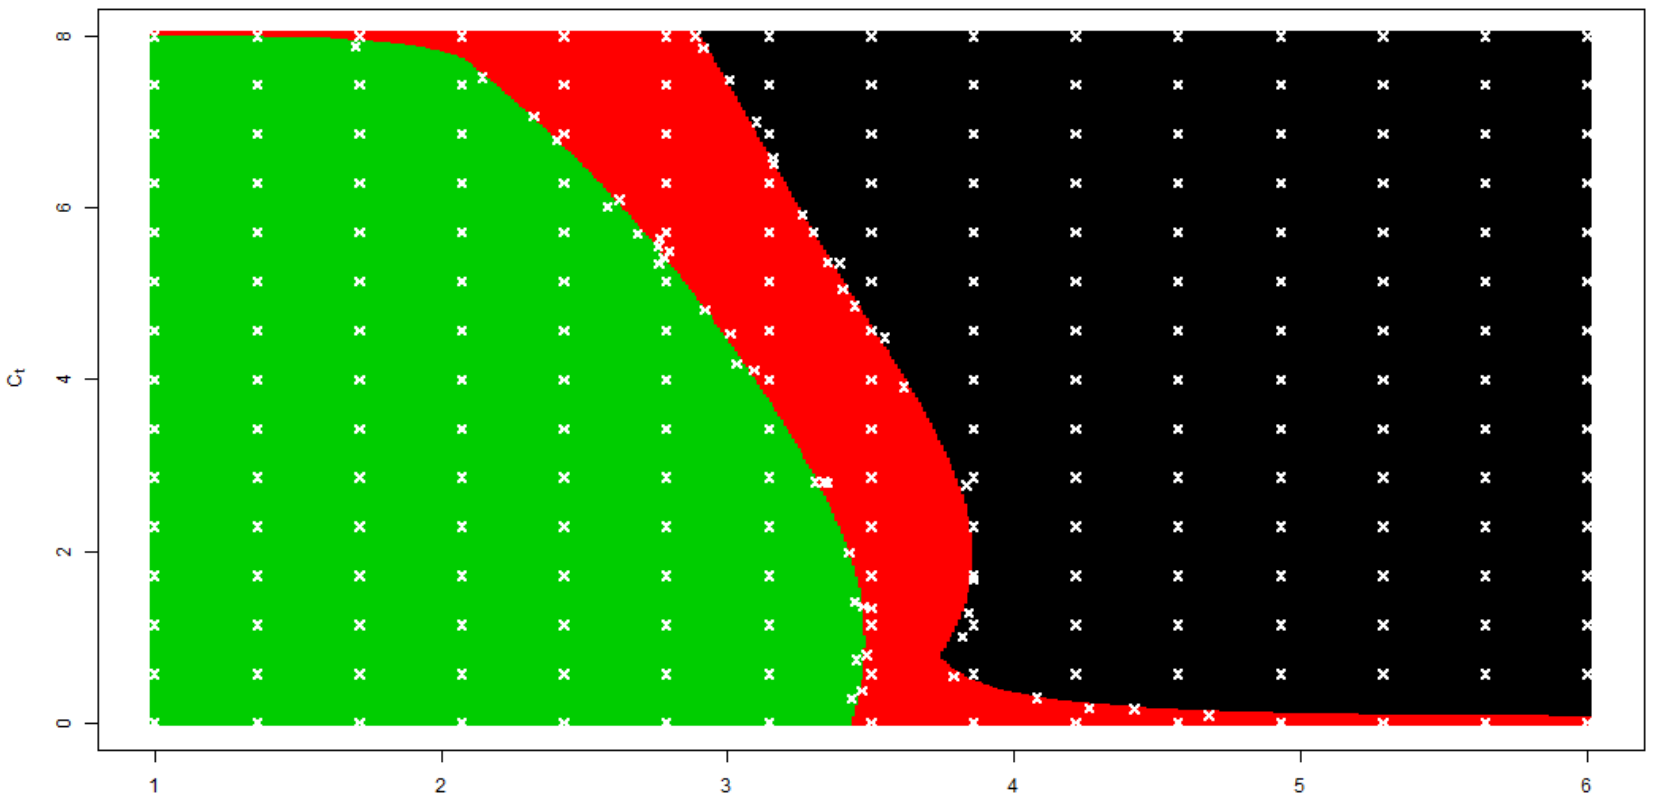

Figure: $\mathcal{C}_{50,store}(x)$,green:inject, red:store, black:withdraw

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Illustration for $t = 50$

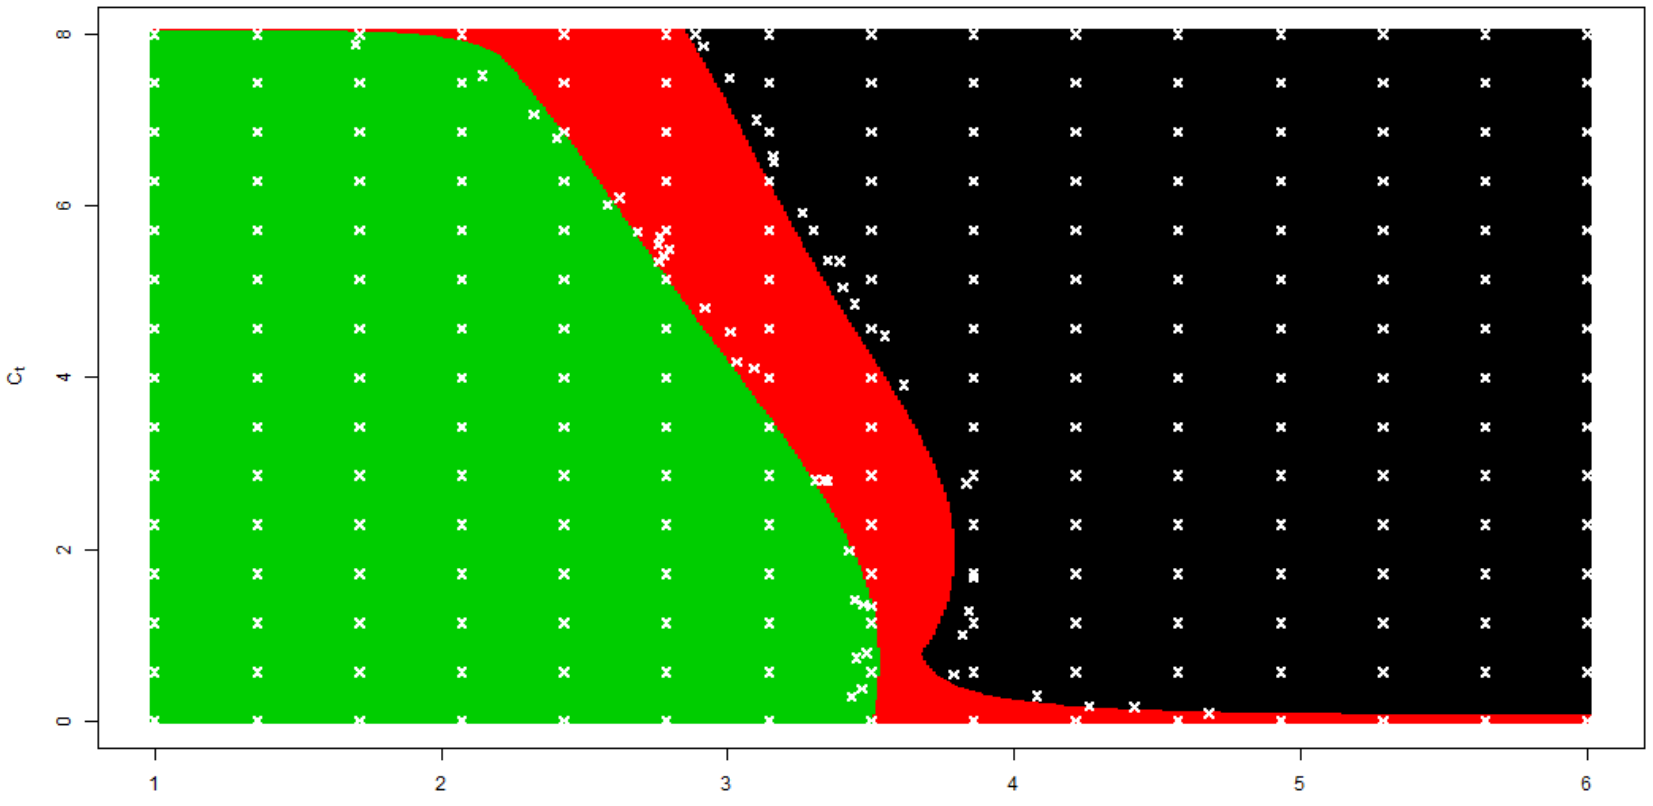

Figure: $\mathcal{C}_{50,store}(x)$,green:inject, red:store, black:withdraw

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Illustration for $t = 50$

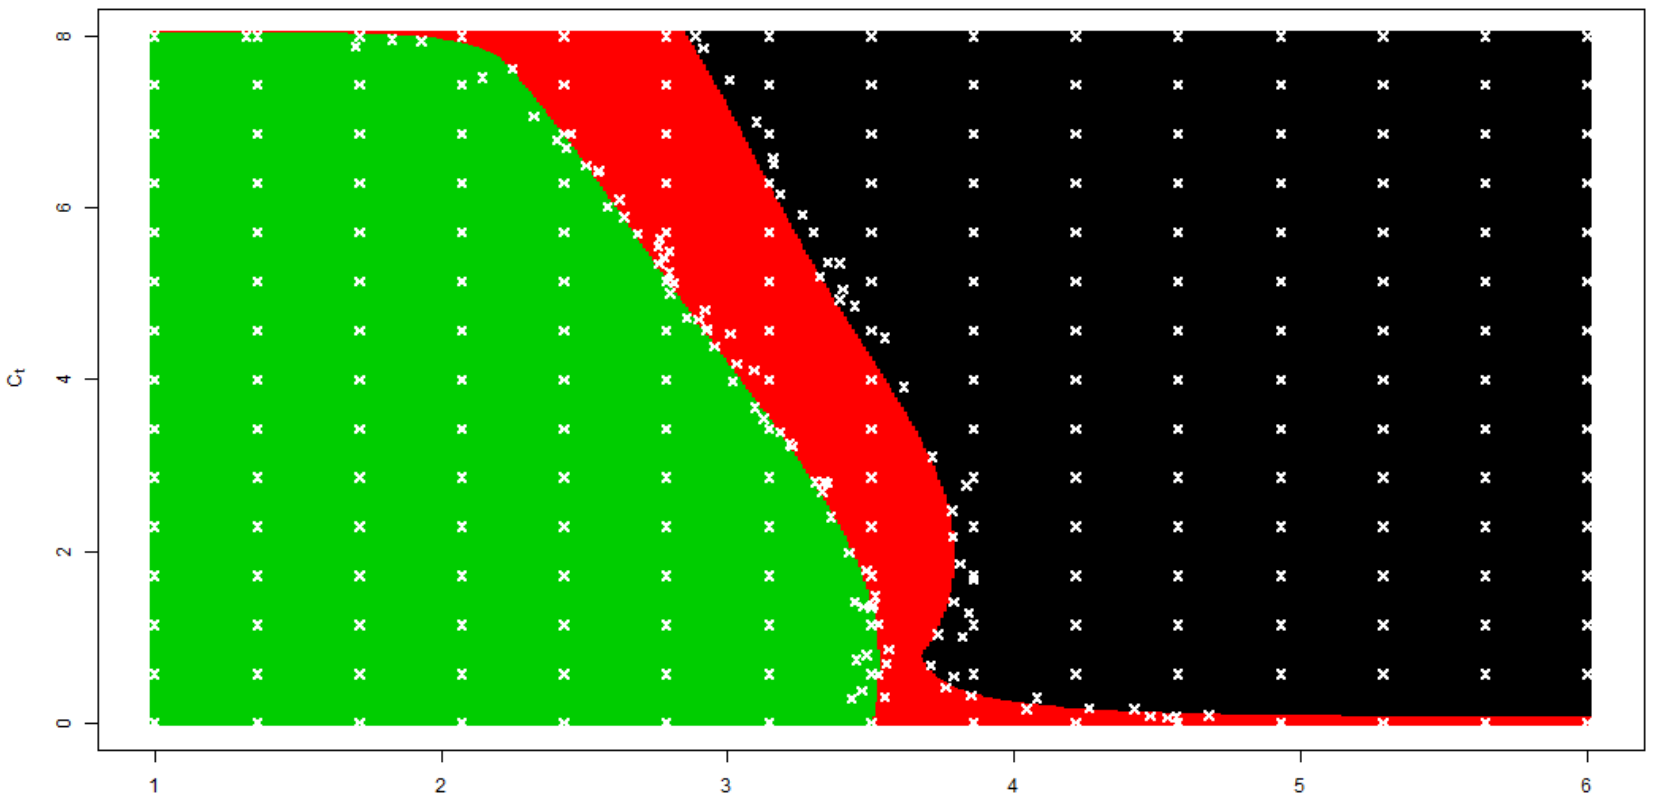

Figure: $\mathcal{C}_{50,store}(x)$,green:inject, red:store, black:withdraw

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Illustration for $t = 50$

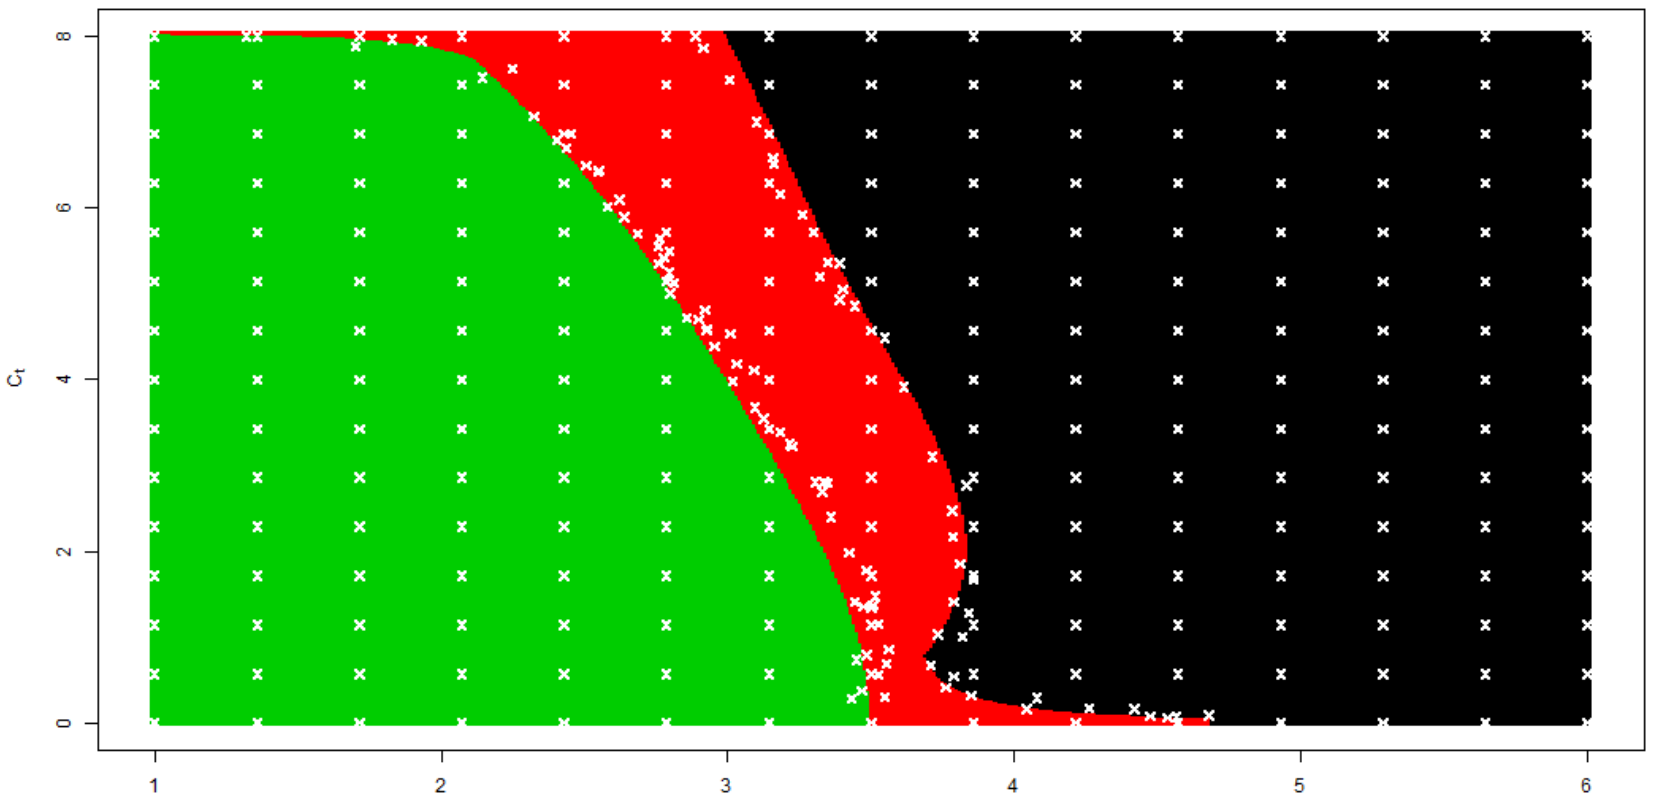

Figure: $\mathcal{C}_{50,store}(x)$,green:inject, red:store, black:withdraw

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Illustration for $t = 50$

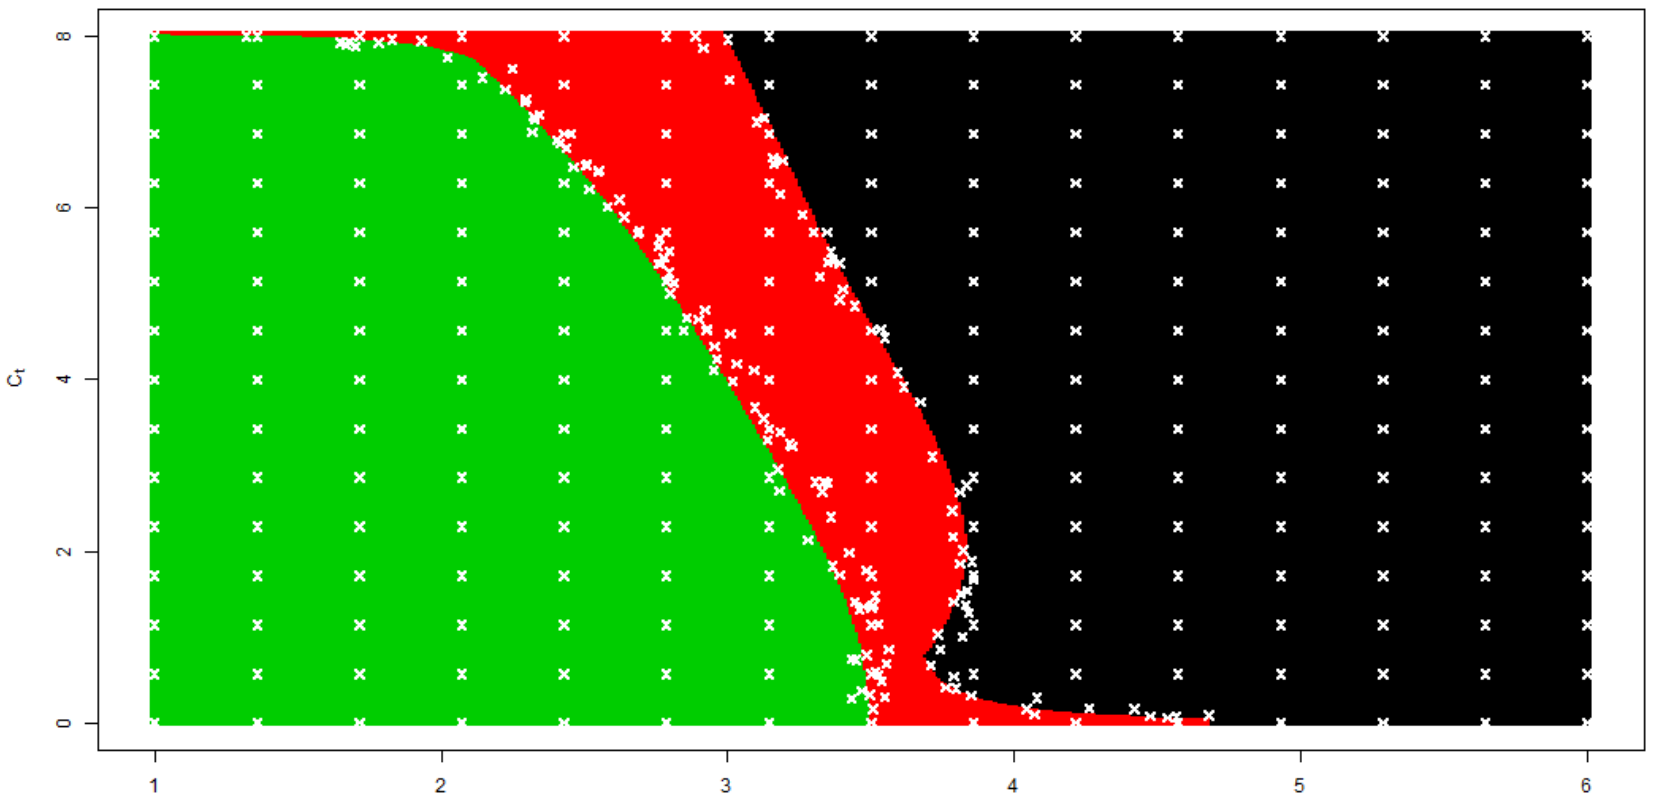

Figure: $\mathcal{C}_{50,store}(x)$,green:inject, red:store, black:withdraw

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Illustration for $t = 50$

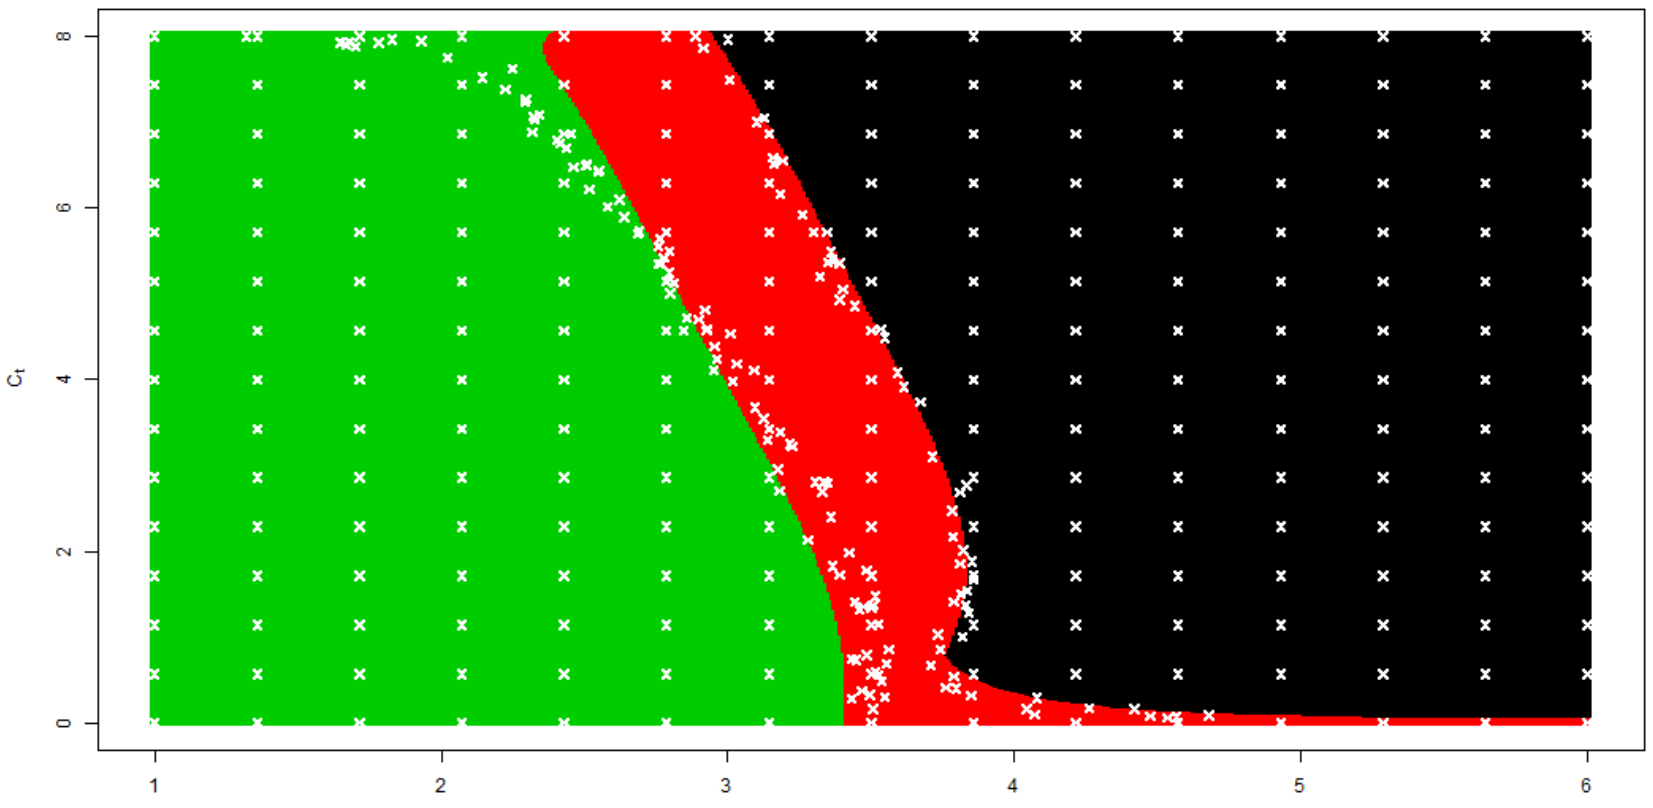

Figure: $\mathcal{C}_{50,store}(x)$,green:inject, red:store, black:withdraw

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Illustration for $t = 50$

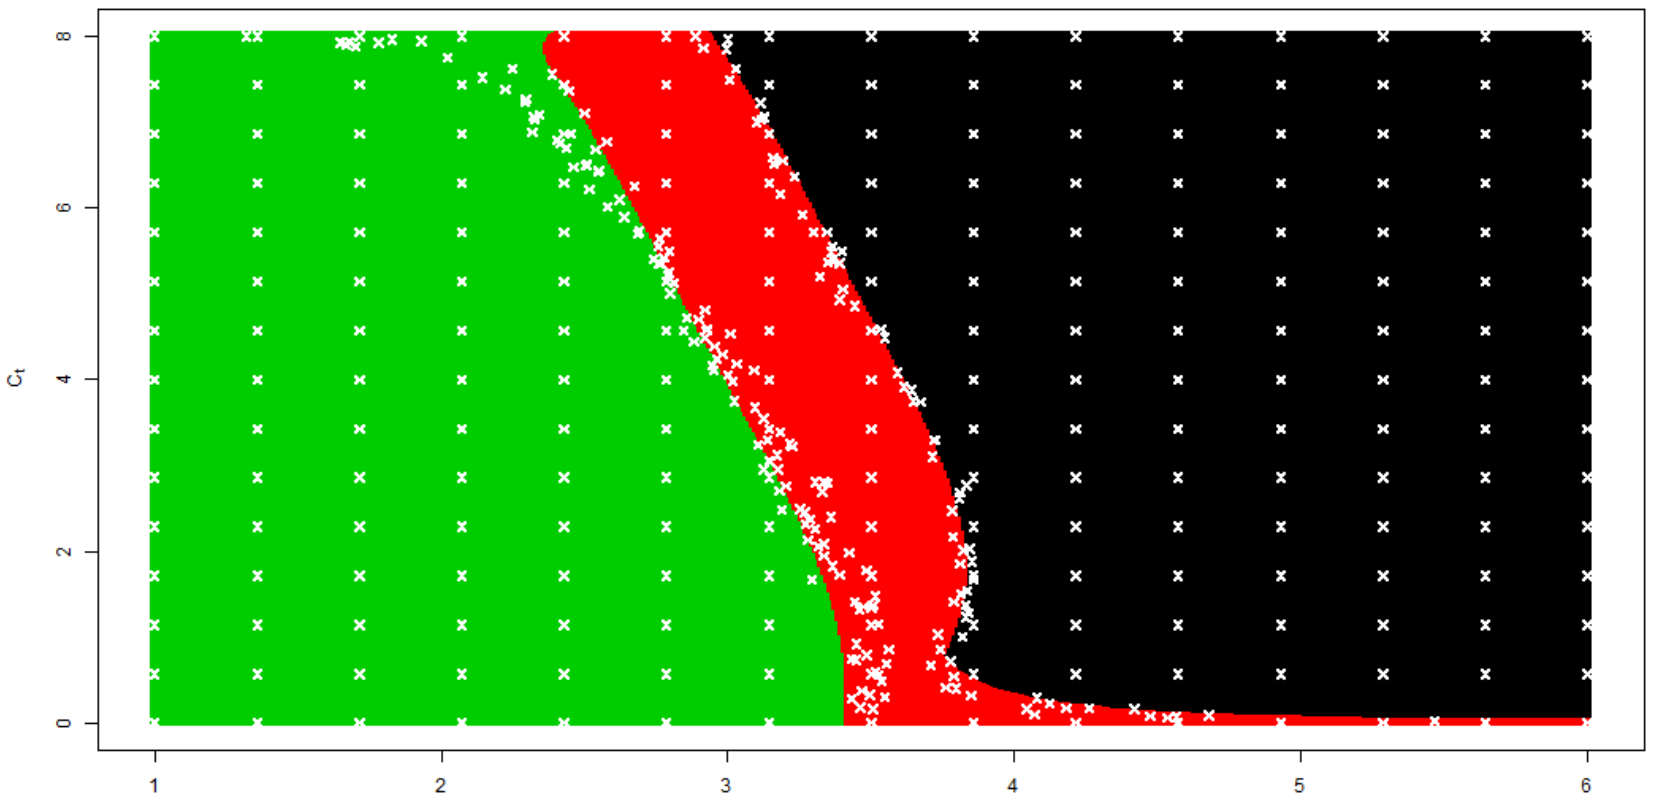

Figure: $\mathcal{C}_{50,store}(x)$,green:inject, red:store, black:withdraw

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Illustration for $t = 50$

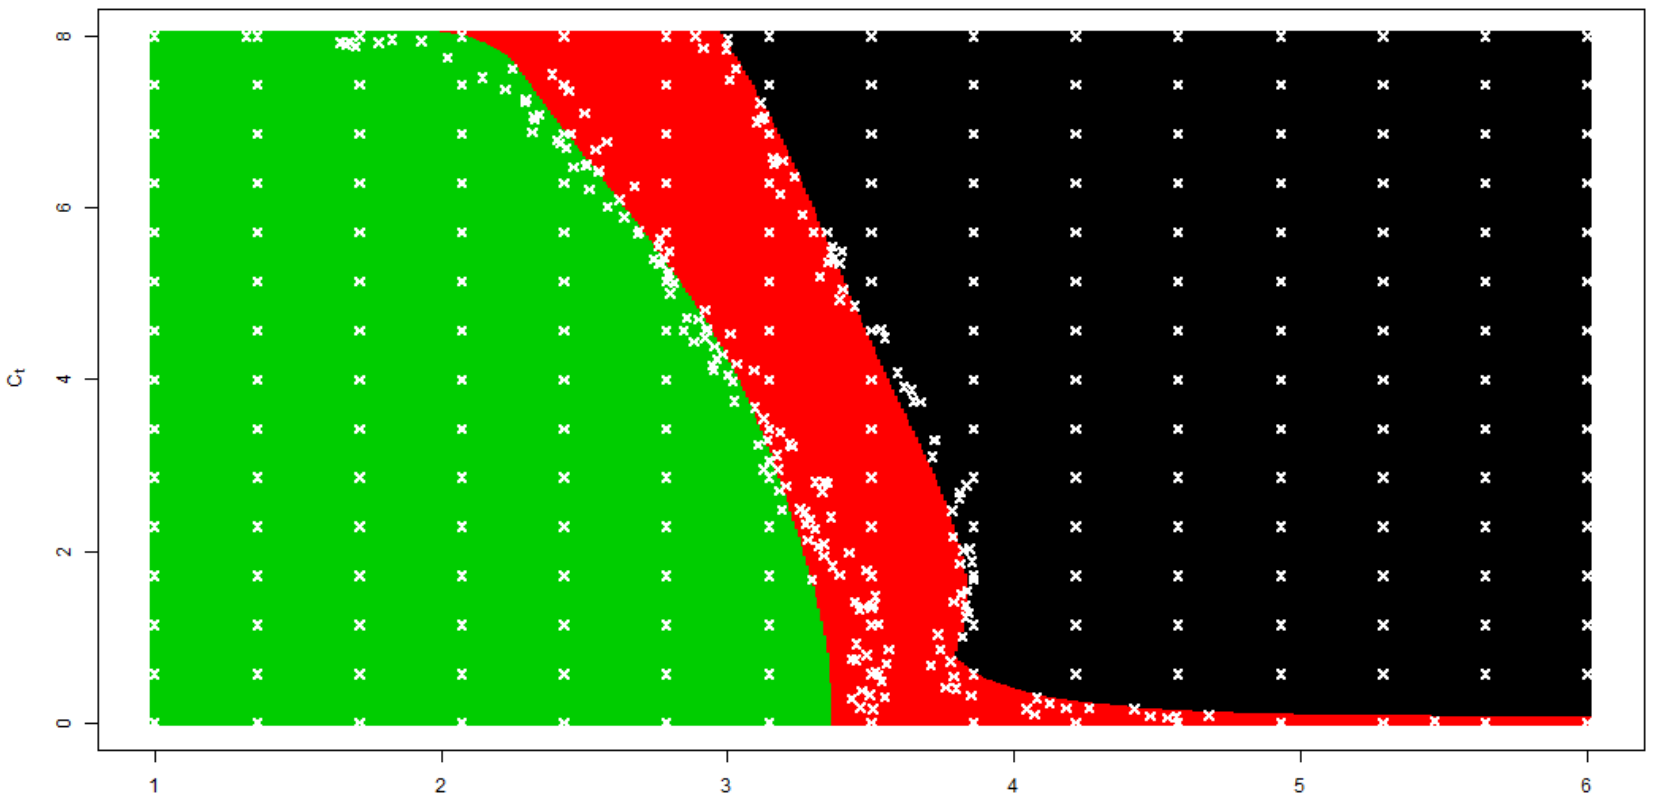

Figure: $\mathcal{C}_{50,store}(x)$,green:inject, red:store, black:withdraw

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Illustration for $t = 50$

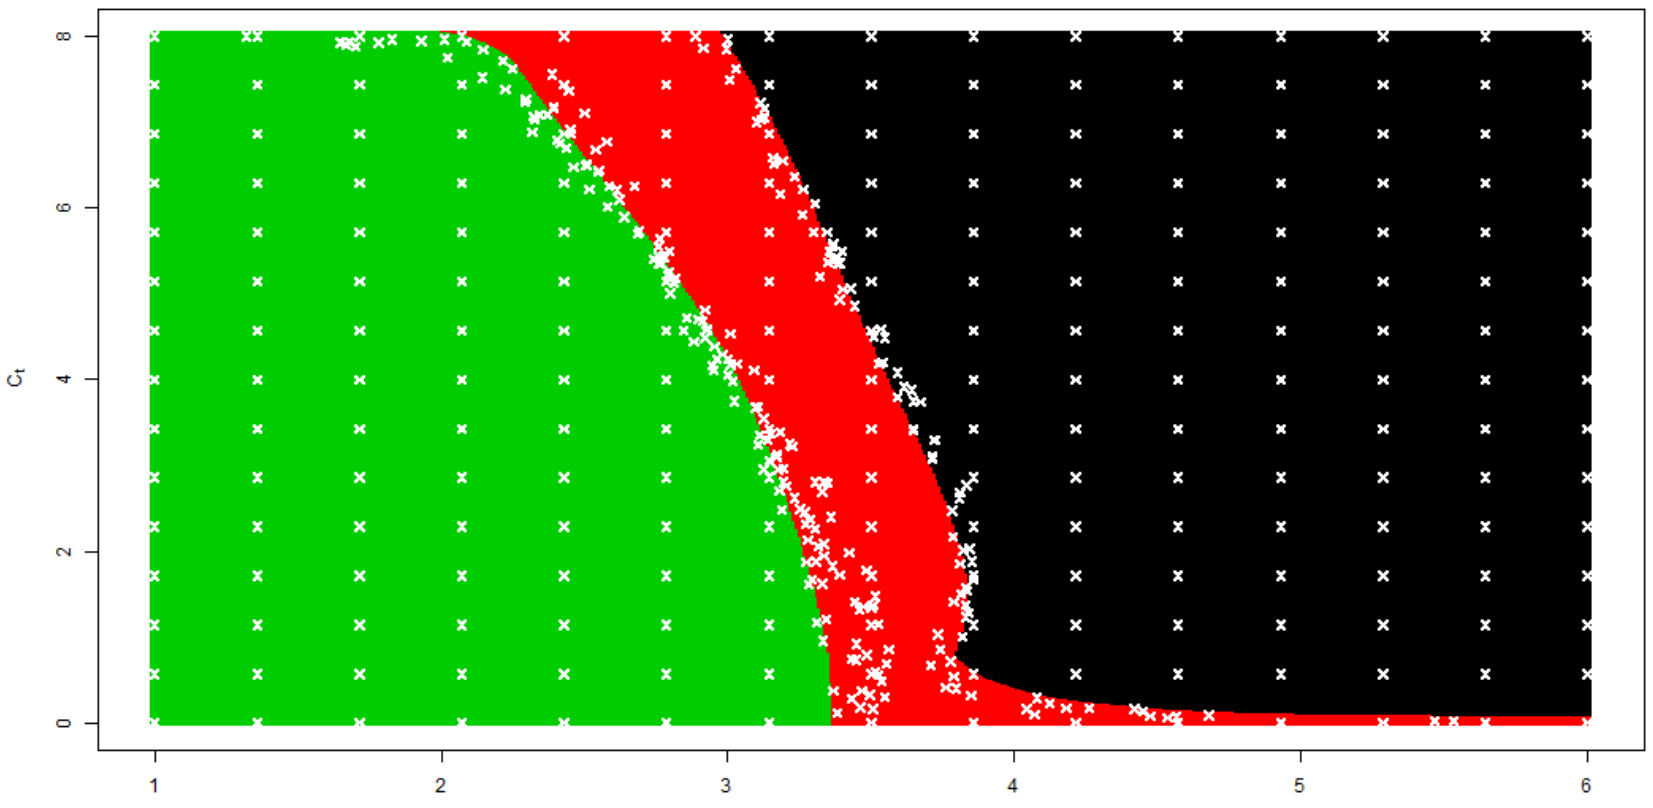

Figure: $\mathcal{C}_{50,store}(x)$,green:inject, red:store, black:withdraw

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Illustration for $t = 50$

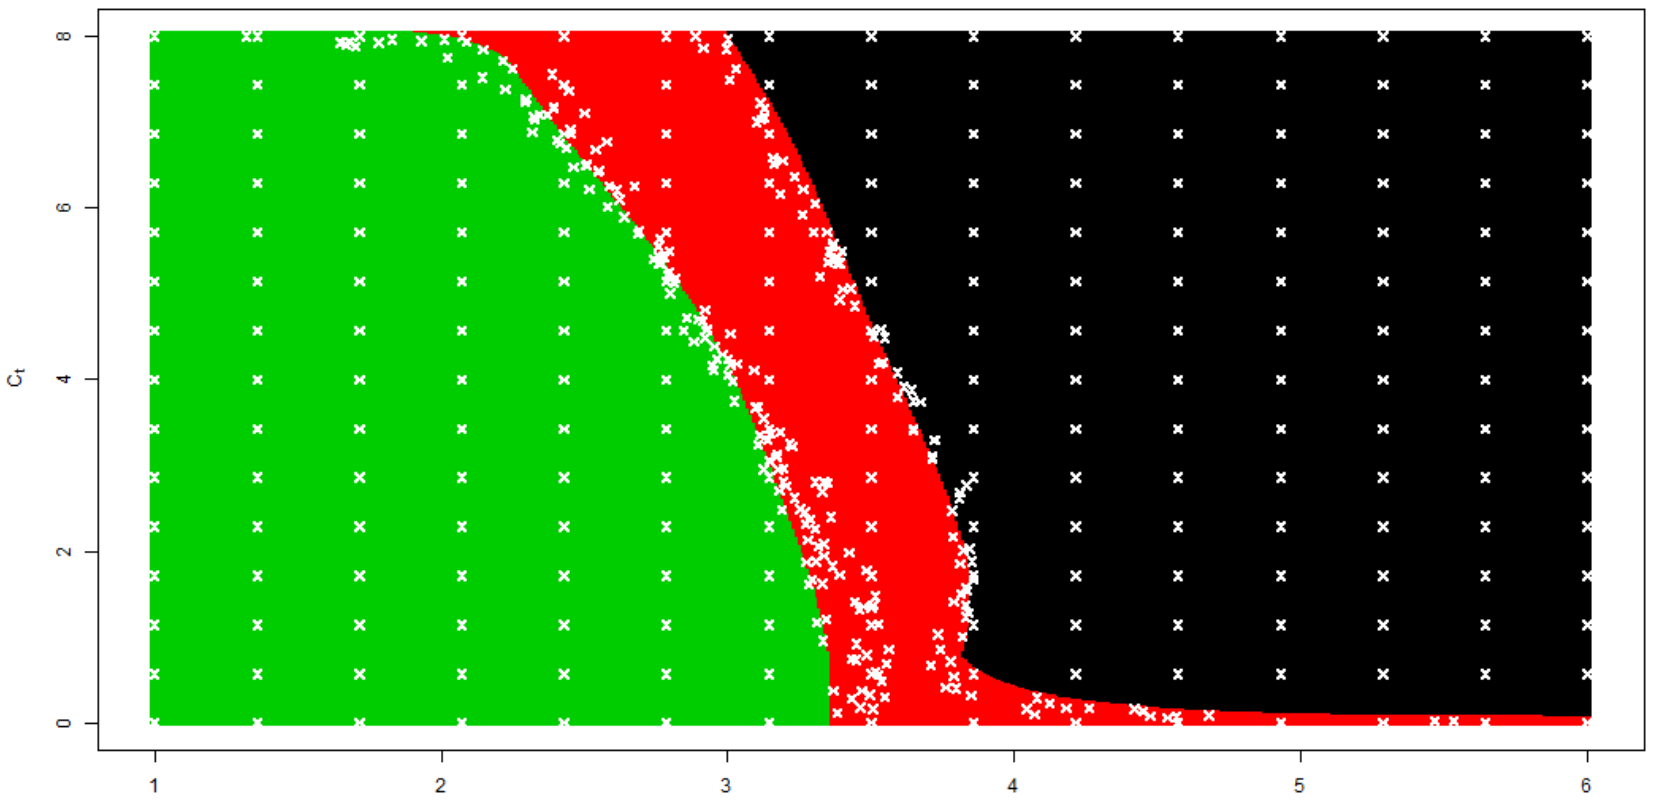

Figure: $\mathcal{C}_{50,store}(x)$,green:inject, red:store, black:withdraw

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Illustration of storage values

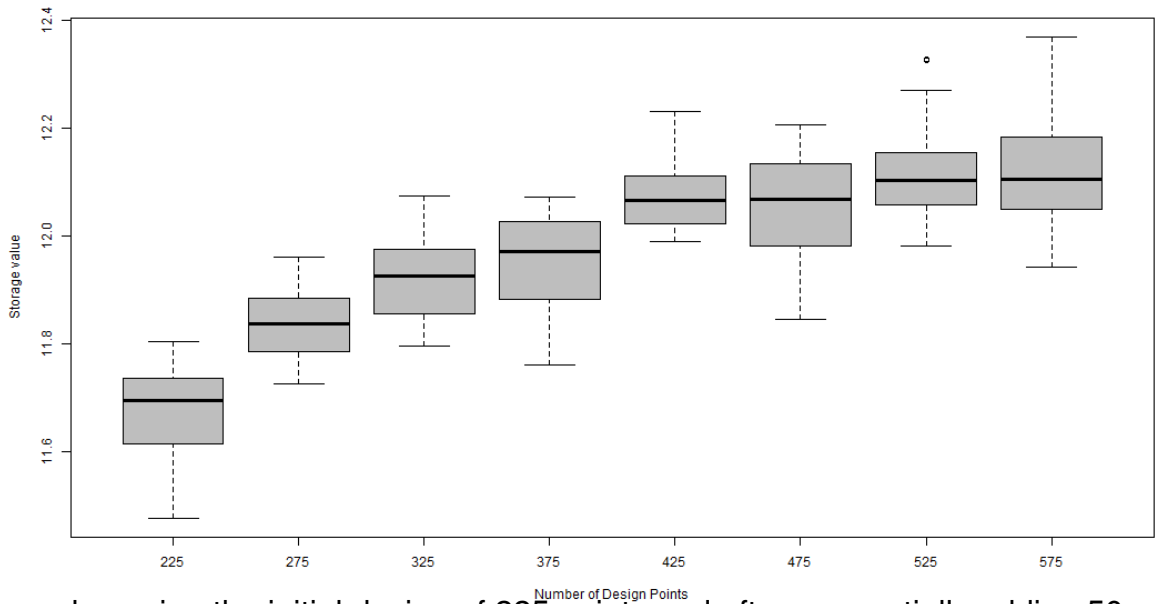

Figure: Storage value using the initial design of 225 points and after sequentially adding 50 points. The values are found by applying the Algorithm on 2500 simulations and calculate the mean of all simulations. The procedure is repeated 20 times at each sampling step.We compare the results with the approach of Carmona and Ludkovski (2010) with 10,000 simulations and polynomial basis functions of order 3 including all possible cross products. Using the same initial states, we find a value of 12.10 MM$, which has a deviation from below 1% to our result.

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Illustration of storage values

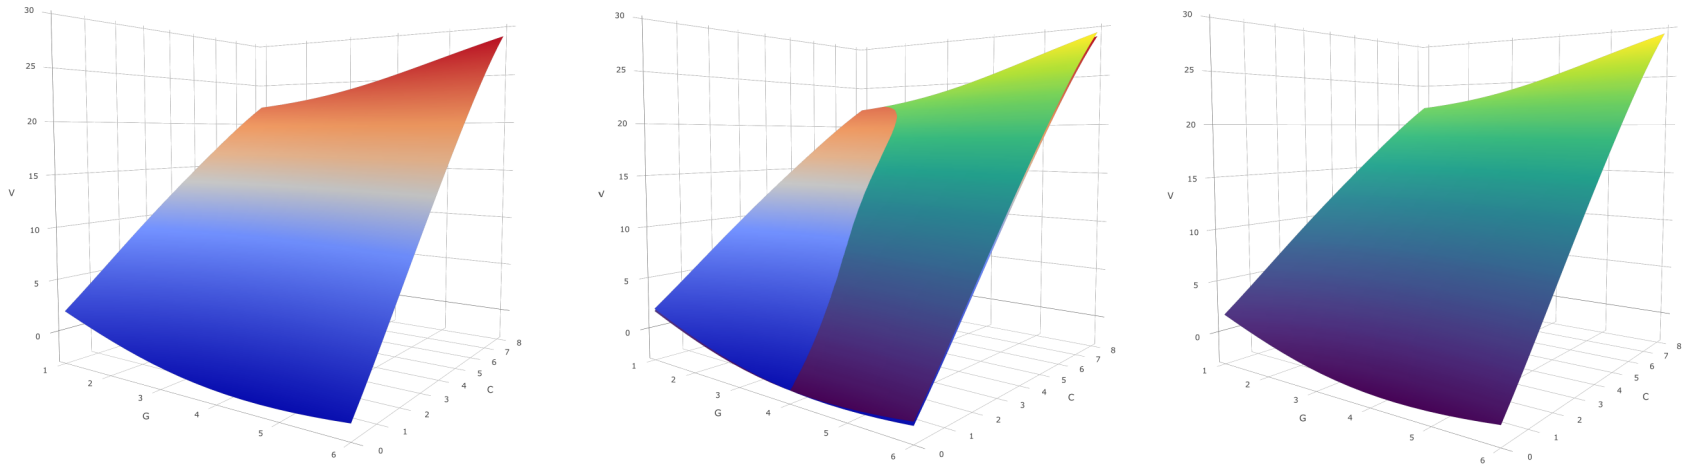

Figure: Posterior mean functions $m_i^{(n)}(x_u)$, which represent the storage value at time $t = 1$ for state $i = \textit{inject}$ (left), $i = \textit{withdrawal}$ (right) and both surfaces interleaved (middle). The contour indicates the states, in which the storage value is equal.

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

Discussion

- General form of the (continuation) value function
- No selection of the basis functions
- On-line learning: refine the solution without loosing the information from the initial optimization if the initial solution is not sufficient
- Potential for parallelization
- nested diagnostic tools in the kriging model
- numerically expensive

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

References I

Balata, A. and Palczewski, J. (2017). Regress-later monte carlo for optimal inventory control with applications in energy. *arXiv preprint arXiv:1703.06461*.

Boogert, A. and De Jong, C. (2008). Gas storage valuation using a monte carlo method. *The journal of derivatives*, 15(3):81–98.

Bouchard, B. and Warin, X. (2012). Monte-carlo valuation of american options: facts and new algorithms to improve existing methods. In *Numerical methods in finance*, pages 215–255. Springer.

Carmona, R. and Ludkovski, M. (2010). Valuation of energy storage: an optimal switching approach. *Quantitative Finance*, 10(4):359–374.

Choudhury, A. R. and King, A. (2008). Optimizations in financial engineering: the least-squares monte carlo method of longstaff and schwartz. In *Parallel and distributed processing, 2008. IPDPS 2008. IEEE International Symposium on*, pages 1–11. IEEE.

Cortazar, G., Gravet, M., and Urzua, J. (2008). The valuation of multidimensional american real options using the lsm simulation method. *Computers & Operations Research*, 35(1):113–129.

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

References II

Gramacy, R. and Ludkovski, M. (2015). Sequential design for optimal stopping problems. *SIAM Journal on Financial Mathematics*, 6(1):748–775.

Hu, R. and Ludkovski, M. (2016). Sequential design for ranking response surfaces. *SIAM/ASA Journal on Uncertainty Quantification (to Appear)*.

Kharroubi, I., Langrené, N., and Pham, H. (2014). A numerical algorithm for fully nonlinear hjb equations: an approach by control randomization. *Monte Carlo Methods and Applications*, 20(2):145–165.

Kiesel, R., Gernhard, J., and Stoll, S.-O. (2010). Valuation of commodity-based swing options. *The Journal of Energy Markets*, 3(3):91.

Kohler, M. (2010). A review on regression-based monte carlo methods for pricing american options. In Devroye, L., Karasözen, B., Kohler, M., and Korn, R., editors, *Recent Developments in Applied Probability and Statistics*, pages 37–58. Physica-Verlag.

Longstaff, F. A. and Schwartz, E. S. (2001). Valuing american options by simulation: A simple least-squares approach. *Review of Financial studies*, 14(1):113–147.

Ludkovski, M. (2016). Kriging metamodels and experimental design for bermudan option pricing. *Journal of Computational Finance (to Appear)*.

House of Energy Markets & Finance

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

References III

Nadarajah, S., Margot, F., and Secomandi, N. (2017). Comparison of least squares monte carlo methods with applications to energy real options. *European Journal of Operational Research*, 256(1):196–204.

Tsitsiklis, J. N. and Van Roy, B. (2001). Regression methods for pricing complex american-style options. *IEEE Transactions on Neural Networks*, 12(4):694–703.