# Bayesian inference on mixtures of Ornstein-Uhlenbeck processes for modelling electricity spot prices, and applications to the economically optimal control of CHP and energy storage

John Moriarty*
Joint work with Jhonny Gonzalez* and Jan Palczewski#

*University of Manchester, #University of Leeds

Heat storage in a flexible district energy system

Electricity price process ‘dimension’ and calibration via MCMC

Results from two electricity markets

# Section 1

## Heat storage in a flexible district energy system

# Toy model of a flexible district energy system

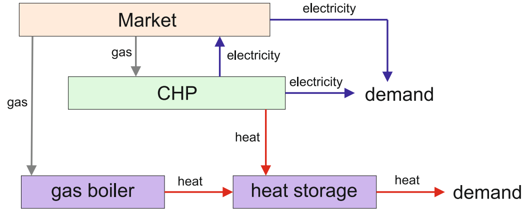


Local electricity and heat demand must be satisfied by market, CHP, boiler and heat store [Kitapbayev, Moriarty and Mancarella, *Applied Energy* 2014]. The heat store:

- Decouples heat supply from heat demand
- Captures excess heat when CHP production is high (eg. during electricity peak demand / electricity price spikes)
- Helps meet heat demand when gas prices are high
- Can act with CHP to provide demand response to price signals in both electricity and gas markets
- Thus can potentially improve the business case for energy flexibility relative to passive load following

# Stochastic optimisation challenge

Under time-varying demand and stochastic prices:

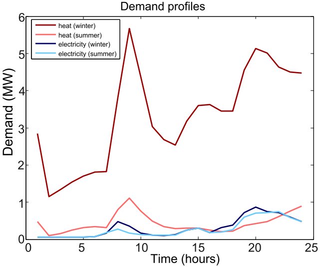


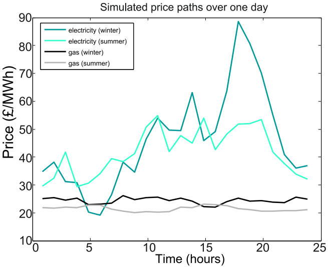


and with flexible operation of CHP and boiler:

**Table 1**
Operating modes for the multi-energy system shown in Fig. 1.

| CHP | | Boiler | |
|---|---|---|---|
| **Regime** | **Input (Gas, MW)** | **Regime** | **Input (Gas, MW)** |
| C1 | 0 | F1 | 0 |
| C2 | 1.875 | F2 | 1.65 |
| C3 | 3.75 | F3 | 3.3 |
| C4 | 5.625 | F4 | 4.95 |
| C5 | 7.5 | F5 | 6.6 |

we wish to minimise costs by optimally choosing the *times* and *states* when switching between operational states *CiFj*.

# Stochastic optimisation of a flexible district energy system

The precise optimisation problem we solve is minimising the expected total net discounted operational cost:

$$V(d) = \min_{u \in \mathbf{u}} \mathrm{E}^{d,u} \left[ \int_t^T e^{-r(s-t)} \psi(u_s) ds \right] \tag{1}$$

- $\mathbf{u}$ is the set of all admissible feedback control policies - so controller is allowed to observe the system state in real time
- $d = (t, g, e, c)$ represents the state of the system at any particular time $t$. The components of $d$ are time, gas price, electricity price and level of stored heat respectively
- $\psi(u(d))$ is the rate of expenditure on both gas and electricity (net of the rate of any electricity sales back to the market) under the particular feedback control policy $u \in \mathbf{u}$ when the system state is $d$

# Numerical method: least squares Monte Carlo regression

- Method based on Carmona and Ludkowski (2010) and references therein
- Learns the stochastic dynamics of the market and energy system (takes 'Monte Carlo' simulations as input)
- Takes account of system constraints and opportunity costs
- Estimates the conditional expectation of the value function $V(d)$ through statistical regression
- Returns estimated optimal feedback controller for real-time demand response
- Also returns the value function $V(d)$ through dynamic programming, for investment analysis (eg. compare with / without heat storage)
- Can be optimised to run in minutes (has been deployed for a UK startup electricity supplier using demand response)

# An example optimal feedback stochastic control policy

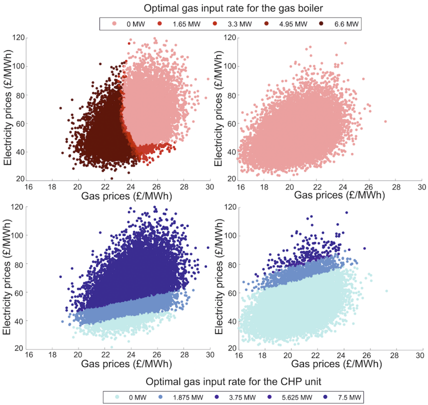


Figure: Optimal feedback controls for boiler (top) and CHP (bottom), in winter (left) and summer (right). Calibrated with UK market data [Kitapbayev et al., 2014]

The feedback strategy is interpretable:

- boiler exploits gas price fluctuations over time
- CHP exploits (instantaneous) spark spread
- boiler acts in sympathy with CHP in winter (not used in summer).

# Considerations on electricity price models

- Electricity prices are spiky and often modelled by jump diffusions
- No previous work on numerical solution of stochastic optimal switching problems driven by jump diffusions
- Clear *a priori*, and suggested by initial numerical experiments, that value functions will be unrealistic if an inappropriate 'dimension' of price processes is used
- So LSM method needs to know **dimension** of the price process...

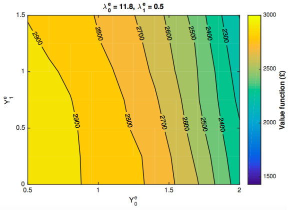


Figure: Contour plot of value function with two-component electricity price model (jump component is vertical and diffusion component is horizontal).

# Section 2

## Electricity price process ‘dimension’ and calibration via MCMC

# Jump-diffusion electricity price models

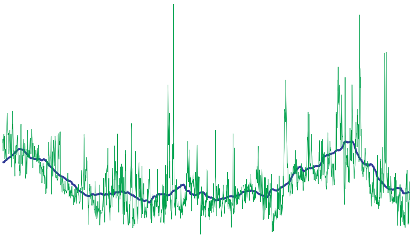

Electricity prices are **spiky** and **mean-reverting** and can be modelled using multiple components, for example as $e^{f(t)}\left(\sum_{i=0}^{n} Y_i(t)\right)$:

- Seasonal component $f$ models weather and consumption / production patterns
- Mean-reverting diffusion component $Y_0$ represents 'normal' price evolution
- Each mean-reverting jump component $(Y_i)_{i\geq 1}$ adds one more 'dimension'

# Price spike modelling

Well calibrated models including spikes are important:

- Price spikes are a risk to buyers but potentially a source of revenue for the CHP unit
- For analysis (eg. LSM method just presented), spikes should be separated from 'normal' price variations
- However spikes decay over *multiple periods*, so:
  - it's not sufficient just to filter out large price movements
  - consecutive jumps mix together producing longer disturbances, making spike identification more challenging

Indicative example:

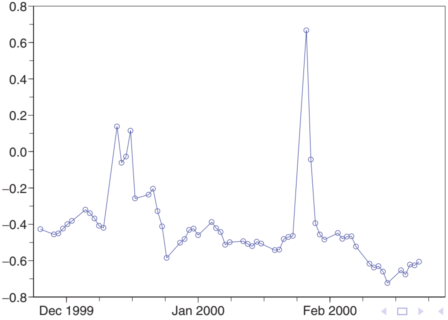

# Price process calibration

The *Markov property* is a key assumption made in stochastic process modelling. Addition of stochastic processes in general destroys the Markov property - so we estimate the individual (Markovian) factors $Y_i$ in the multifactor model $e^{f(t)}\left(\sum_{i=0}^{n} Y_i(t)\right)$.

Meyer-brandis and Tankov (2008) assume a single spike path $Y_1$ (ie. $n = 1$) which can be known with certainty, and propose two non-parametric methods to filter it out (NB: *path*, not just set of jumps). Separate parameter estimates may then be made for the diffusion and jump components.

Taking these parameter estimates as prior information, we seek a fully Bayesian approach making minimal assumptions - to 'let the data speak'. Large number of interdependent parameters to calibrate $\Rightarrow$ try MCMC.

# Aim: Bayesian joint estimation via MCMC

Stochastic model:

$$\text{Spot price } S(t) \quad = \quad e^{f(t)} \left( \sum_{i=0}^{n} Y_i(t) \right),$$

where the 'dimensions' $Y_i$ have mean-reverting stochastic dynamics:

$$dY_i(t) \quad = \quad \lambda_i^{-1}(\mu_i - Y_i(t))dt + \sigma_i dL_i(t), \quad Y_i(0) = y_i.$$

- $\lambda_i$: mean reversion parameters
- $i = 0$: diffusion component, $L_0(t) = W(t)$ is Brownian motion.
- $i \geq 1$: jump components, $L_i$ compound Poisson process with rate $\eta_i > 0$, exponential jump sizes with mean $\beta_i$. So $Y_i, i \geq 1$, are Gamma OU processes.

# Augmented state space and parametrisation

Each jump path, eg. $Y_1$ is completely known given $\Phi_i = \{(\tau_{j,}\xi_j)\}$ its set of arrival times and corresponding jump sizes of $L(t)$, $\lambda_1$ and $Y(0)$, since

$$Y_i(t) = Y_1(s)e^{-\lambda_1^{-1}(t-s)} + \sum_{s<\tau_j\leq t} e^{-\lambda_1(t-\tau_j)}\xi_j, \qquad t > s.$$

We 'add dimensions' to the observed price by employing a *latent variable* model which augments the observed price process $X$ with these states $\Phi_i$ (this parametrisation gives good mixing properties for the MCMC procedure).

# MCMC updates

MCMC 'fills in' the missing dimensions by a sophisticated and delicate form of trial and error. The MCMC procedure updates $\Phi$ with random combinations of:

- Birth and death proposal: Place a new jump with probability $p$, kill an existing jump with probability $p-1$.
- Local displacement proposal:
  - Choose randomly one of the jump times, say, $\tau_j$, and generate a new jump time $\tau_{new}$ uniformly on $[\tau_{j-1}, \tau_{j+1}]$ (and deterministically re-size this jump for consistency)
- Block update of all jump sizes (proposal variance is inversely proportional to current number of jumps).

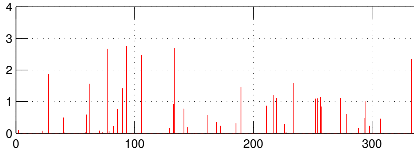

# Does it work? Simulation efficiency with three factors

- Red: Two independent simulated jump processes. Added together with a simulated diffusion (not shown) and MCMC procedure applied to the sum.
- Blue: Representation of the posterior distribution of jump components.

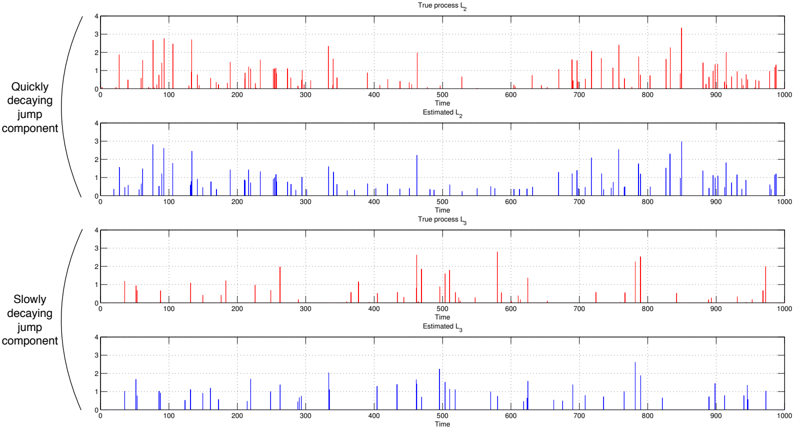

Section 3

# Results from two electricity markets

# Results from two electricity markets

- The least squares Monte Carlo optimisation is capable of learning the stochastic dynamics of a multifactor electricity price
- Since dynamic programming is used, the price process must be *Markovian* - using an inappropriate number of components will violate this key assumption of the optimisation
- We therefore aim to demonstrate that the MCMC procedure can find the appropriate number of Markovian components.
- **Minimal criterion for success**: we want $Y_0$ to look like a diffusion, ie. the increments of the fitted Browian motion $L_0$ should 'pass' the one-sample Kolmogorov-Smirnov test. (A Bayesian posterior $p$-value may be obtained by averaging the $p$-values over the MCMC posterior; we seek $p > 0.1$)

We examine two different electricity spot markets:

- daily APXUK data from March 27, 2001 to November 21, 2006
- daily EEX data from June 16, 2000 to November 21, 2006

with weekends removed.

# Results from two electricity markets

APXUK data:

- The two factor (ie. one diffusion factor, one jump factor) model implies a diffusion process with posterior mean $p$ value $=0.06$
- With three factors (ie. one diffusion, two jump factors), the posterior mean $p$ value increases to 0.33. The MCMC fits an additional, slowly decaying jump component with mean reversion rate $\approxeq 5$ times lower

EEX data:

- Logarithmic prices used to better remove seasonality
- The two factor model implies a diffusion with posterior mean $p$ value $=0.002$
- Using three factors as above (ie. $Y_0(t) + Y_1(t) + Y_2(t)$) is not helpful this time, as the MCMC procedure apparently fails to converge
- Subtracting (rather than adding) the third factor (ie. $Y_0(t) + Y_1(t) - Y_2(t)$) yields MCMC convergence, with a posterior mean $p$ value of 0.23. So the MCMC successfully fits negative jumps to the log prices.

# Posterior for jump process (APXUK, two factors)

Representation of posterior for a single jump process, and the implied diffusion process (posterior $p = 0.06$):

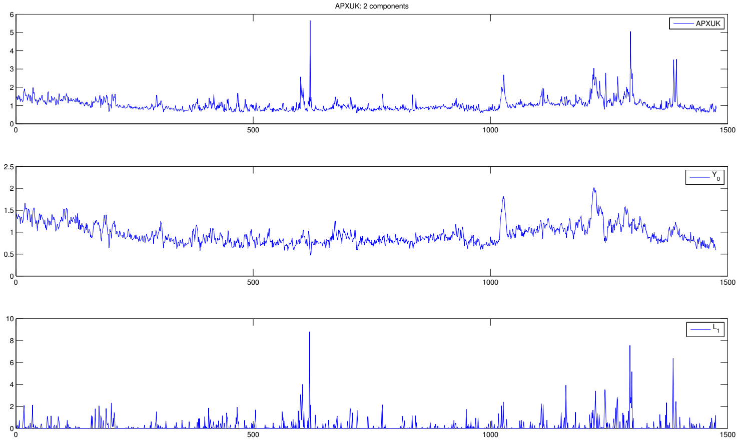

# Posterior for jump process (APXUK, three factors)

Representation of the posterior for two jump processes, and implied diffusion process (posterior $p = 0.33$):

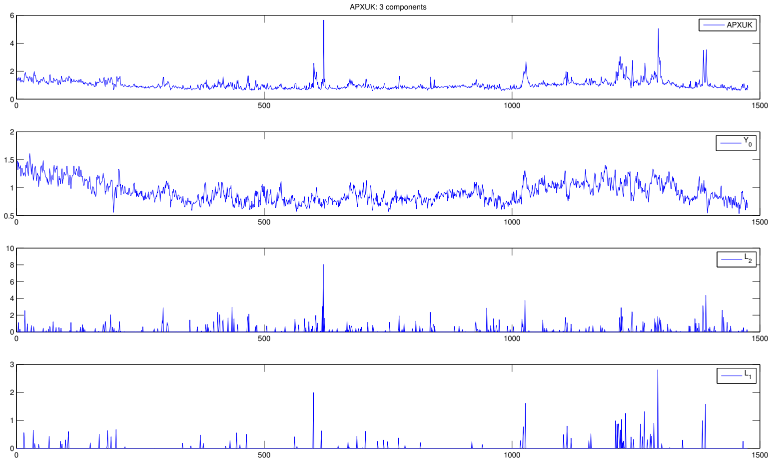

# Posterior for jump process (EEX, two factors)

Representation of the posterior for a single jump process, and implied diffusion process (posterior $p = 0.02$):

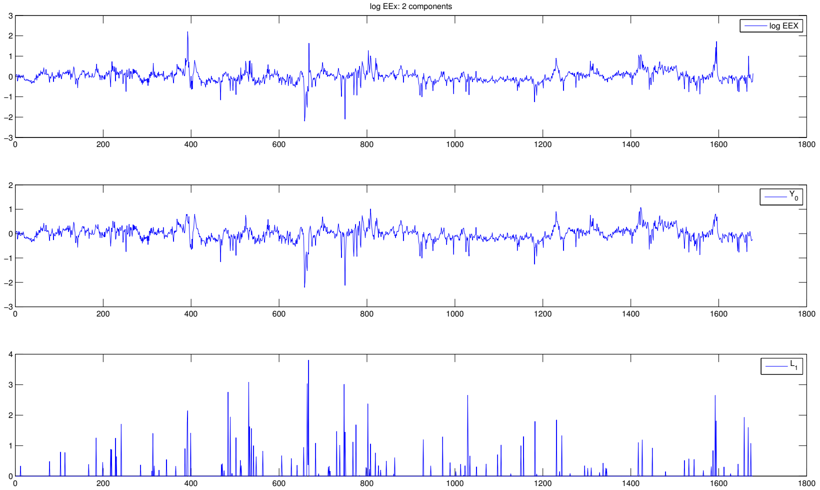

# Posterior for jump process (EEX, three factors)

Representation of the posterior for the jump processes, and implied diffusion process (posterior $p = 0.23$):

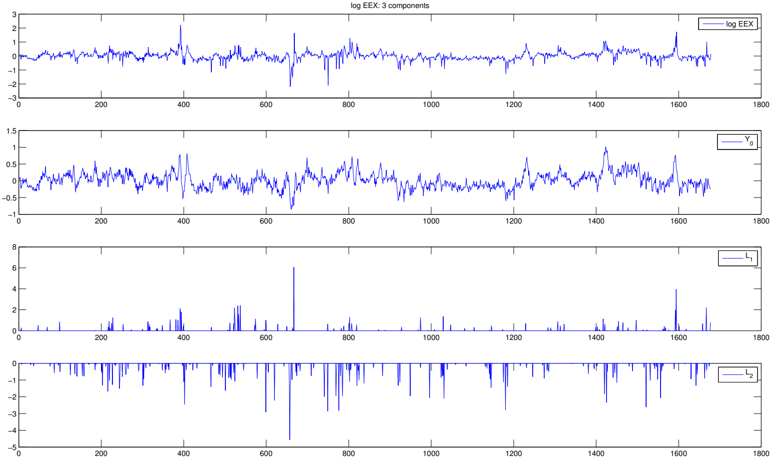

# Conclusions

- We have developed a Bayesian approach to model calibration for multifactor jump-diffusion electricity price models via MCMC, to 'let the data speak'
- Applied to two different electricity spot markets (APXUK and EEX)
- Two jump factors found to be necessary in both cases, but for different reasons: either slowly decaying jumps (APXUK) or negative jumps (log EEX) were required
- Identifying the appropriate number of jump components in electricity prices allows proper application of Markovian numerical optimisation procedures for flexible energy systems.