By Reza Pourmoayed and Lars Relund Nielsen

Culling pigs under price fluctuations

In the production of finishing pigs, pig marketing refers to a sequence of culling decisions until the production unit is empty. The profit of a production unit is highly dependent on the price of pork, the cost of feeding and the cost of buying piglets. Price fluctuations in the market consequently influence the profit, and the optimal marketing decisions may change under different price conditions. We formulate a stochastic dynamic optimization model that optimize the expected reward per time unit. The state of the system is based on information about pork, piglet and feed prices. Moreover, the information is updated using a Bayesian approach and embedded into the model.

Pig production in Denmark

In Denmark approximately 30 million piglets are produced in every year. A significant number of these piglets are exported to other countries (around 11 mill) and the rest of them (approx. 19 million) are sent to fattening units in Denmark. The number of pig farms in Denmark are approximately 3600 where 50% are finishing farms, 30% are integrated farms (both sow and finishing pigs), and the remaining 20% are sow farms (Danish Agriculture and Food Council, 2015). Pork constitutes about 5% of the Danish export representing a profit of about 30 billion DKK per year (Landbrug & Fødevarer, 2015).

In the production of finishing pigs, i.e. from inserting the piglets (with a weight of approx. 30 kg) into the finishing unit until marketing/culling the pigs for slaughter (with a weight of approx. 100-110 kg), one of the important operational decisions is *marketing* of pigs for slaughter at pen level. It refers to a sequence of culling decisions until the production unit is empty.

In the finishing unit, animals are grouped at different levels: herd, section, pen, and animal. Herd is a group of sections, a section includes some pens, and a finisher pen involves some animals (usually 15-20).

In general pigs grow with different growth rates in the pen and obtain their slaughter weight at different times. Hence, during the last weeks of the growing period, the decision maker should determine which pigs should be selected for slaughter (*individual marketing*). Next, after a sequence of individual marketings, the decision maker should decide when to terminate the whole pen, i.e. all the remaining pigs are sent to the abattoir (pen termination) and the pen is prepared for a new batch of pigs. For more information about the production process of finisher pigs see Pourmoayed and Nielsen (2014).

Decisions when considering the finishing unit must be taken in a stochastic production environment. First, pigs do not grow with the same growth rate, i.e. there will be a high degree of uncertainty about the weight of the pigs during the growing period. Moreover, the reward of marketing a pig depends on the *pork price* of the carcass weight, the cost of buying the piglet, i.e. the *piglet price*, and the cost of feeding (*feed price*) at the time when the feed stock is bought. That is, weekly fluctuations in pork, feed, and piglet prices may have an impact on decisions on when to market the pigs and buy more feed. Hence, it is relevant to take into account stochastic elements when modelling marketing decisions.

We model sequential marketing decisions under price fluctuations at pen level using a two-level *hierarchical Markov decision process (HMDP)*. The state of the system is based on information about pork, piglet and feed prices. The model considers time series of pork, piglet and feed prices obtained from the market and a learning approach based on *Bayesian updating* is applied to update price information using the historical data which is embedded into the HMDP. More precisely, state space models for Bayesian forecasting (West and Harrison, 1997) are employed to update the future estimates of pork, piglet and feed prices on a weekly basis. For more details about the Bayesian forecasting of price information, the interested reader can refer to Pourmoayed (2016, Sec. 3.4).

Modelling approach

Marketing decisions are modelled using a two-level HMDP. An HMDP is an extension of a *semi-Markov decision process (semi-MDP)* (Tijms, 2003, Chapter 7) where a series of finite-horizon semi-MDPs are combined into one infinite time-horizon process to the founder level called the founder process (Kristensen and Jørgensen, 2000). The idea is to expand stages of a process to so-called child processes, which again may expand stages further to new child processes leading to multiple levels.

Figure 1 illustrates an HMDP with two levels using a *state-expanded hypergraph* (Nielsen and Kristensen, 2006). At the first level, a single infinite-horizon *founder process* p^0 is defined. Let p^{l+1} denote a *child process* at level $l+1$. Process p^{l+1} is uniquely defined by a given stage n^l, state i^l and action a^l of its *parent process* p^l. For instance, the semi-MDP p^1 in Figure 1 is defined at stage n^0, state i^0 and action a^0 of the founder process p^0 symbolized by the notation $p^1 = (p^0 \,||\, (n^0, i^0, a^0))$. At the lowest level (Level 2 in Figure 1) the HMDP consists of a set of finite-horizon semi-MDPs.

A policy is a decision rule/function that assigns to each state in a process a (jump) action. That is, choosing a policy corresponds to choosing a single hyperarc out of each node in Figure 1. Given a policy, the reward at a stage of a parent process equals the total expected rewards of the corresponding child process. For instance, in Figure 1, the reward of choosing action a^0 in state i^0 at stage n^0 in process p^0 equals the total expected reward of process p^1. Different optimality criteria may be considered. In this paper, our optimality criterion is to maximize the *expected reward per time unit* and the optimal policy of the HMDP is found using a modified policy iteration algorithm. For a detailed description of the algorithm, the interested reader may consult Kristensen and Jørgensen (2000) and Nielsen and Kristensen (2014).

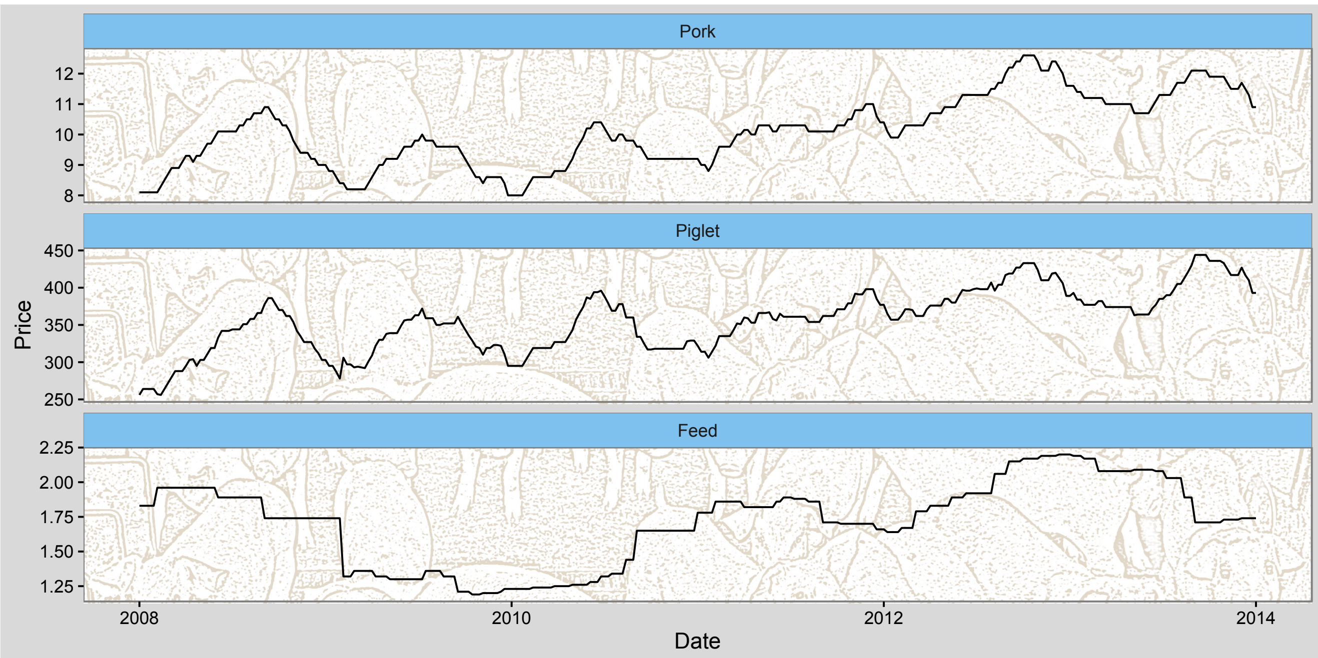

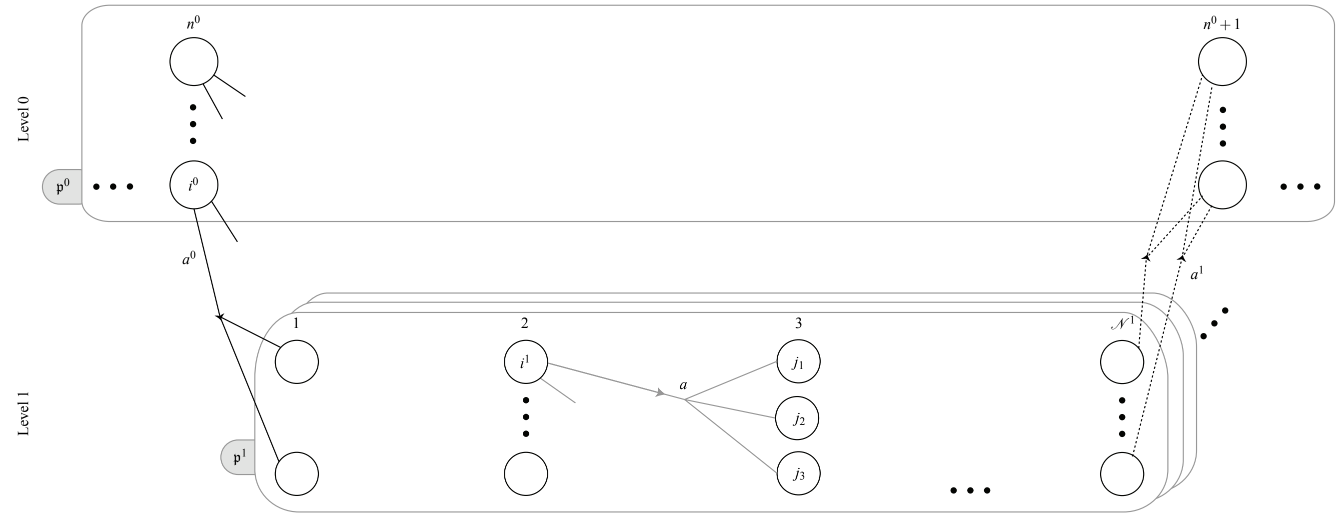

Figure 1: An illustration of a stage in an HMDP (Pourmoayed, 2016). At the founder level (Level 0) there is a single infinite-horizon founder process p^0. A child process, such as p^1 at Level 1 (oval box), is uniquely defined by a given stage, state (node), and action (hyperarc) of its parent process and linked with the parent process using its initial probability distribution (solid lines) and its terminating actions (dashed lines). Each process at level 2 is a semi-MDP. Note that only a subset of the actions have been shown in the figure.

Model

We consider a pen with maximum 18 pigs. The piglets are inserted into the pen in the beginning of week 1 and are growing there for maximum 15 weeks. Since pigs in general grow with different growth rates, they obtain their slaughter weight at different times and hence during the last weeks of the growing period (from week 9 to 15) the farmer should determine which pigs should be selected for slaughter (individual marketing). Theses decisions are taken on a weekly basis and a decision must be taken some days before each delivery. Next, after a sequence of individual marketings, the farmer must decide when to terminate the pen. Terminating a pen means that the remaining pigs in the pen are sent to the slaughterhouse (in one delivery) and after cleaning the pen, another group of piglets (each weighing approx. 30 kg) is inserted into the pen and the *production cycle* is repeated.

At the start of a new production cycle, a new batch of piglets and the required feed stock are bought using the market piglet and feed prices and, within the marketing window (week 9 to 15), pigs are sold to the abattoir using a *pork price function*.

In order to model sequential marketing decisions under price fluctuations using the two-level HMDP, stages, states, actions, rewards and the transition probabilities of each process must be defined. We give a rough description below. For further details see Pourmoayed (2016, Sec. 3).

Founder process p^0

• Stage: A production cycle of 18 pigs, i.e. from inserting the piglets into the pen until terminating the pen.

• Time horizon: Infinite (the pen is filled and emptied an infinite number of times).

• States: A state i^0 represents our information about the market pork, piglet and feed prices.

• Actions: A single jump action a^0 representing insertion of the piglets into the pen.

• Rewards: The cost of buying new piglets at the start of the production cycle.

Child process $p^1 = (p^0 \,||\, (n^0, i^0, a^0))$

• Stage: The first stage represents the period from insertion of the piglets (week 1) until the start of marketing decisions (week 9). The remaining stages represent a week in the marketing period (weeks 9 to 15).

• Time horizon: The maximum number of stages is 15-9+2.

• States: Given a stage a state is defined using two state variables representing:

  • Information related to price deviations from the pork, piglet and feed price given in state i^0, acquired using Bayesian updating.
  • Number of pigs in the pen.

• Actions: At the first stage, marketing is not possible and the production process continues.

At the next stages, the possible actions are “continue”, “terminate”, and a set of actions implying that the *q* heaviest pigs are culled (individual marketing). Finally, at the last stage the pen must be terminated.

• Rewards: Given a state, the expected rewards for actions “continue”, “terminate” and “cull *q* pigs” are calculated using the expected revenue from selling the pigs minus the expected cost of feeding the remaining pigs.

The transition probabilities of the model depends on the statistical model used for Bayesian updating of the price information (for more details see Pourmoayed (2016, Sec 3.4)).

Results

To find the optimal policy of the HMDP, the model was coded using the C++ and R (R Core Team, 2015). The optimal policy was calculated using the modified policy iteration algorithm using the R package ”MDP” (Nielsen, 2009). The source code is available on-line (Pourmoayed and Nielsen, 2015).

To see the behaviour of the optimal policy under different patterns of price fluctuations we consider three scenarios, illustrated in Figure 2, over a period of 15 weeks assuming that the production cycle starts at week one and ends at the start of week 15 at the latest:

Scenario 1: Favourable trend of pork price and unfavourable trends of feed and piglet prices. Pork price increases from 10.3 to 11.3 DKK, feed price increases from 1.79 to 1.92 DKK and piglet price increases from 336 to 396 DKK. This scenario is based on the historical data from weeks 11-25 in 2012.

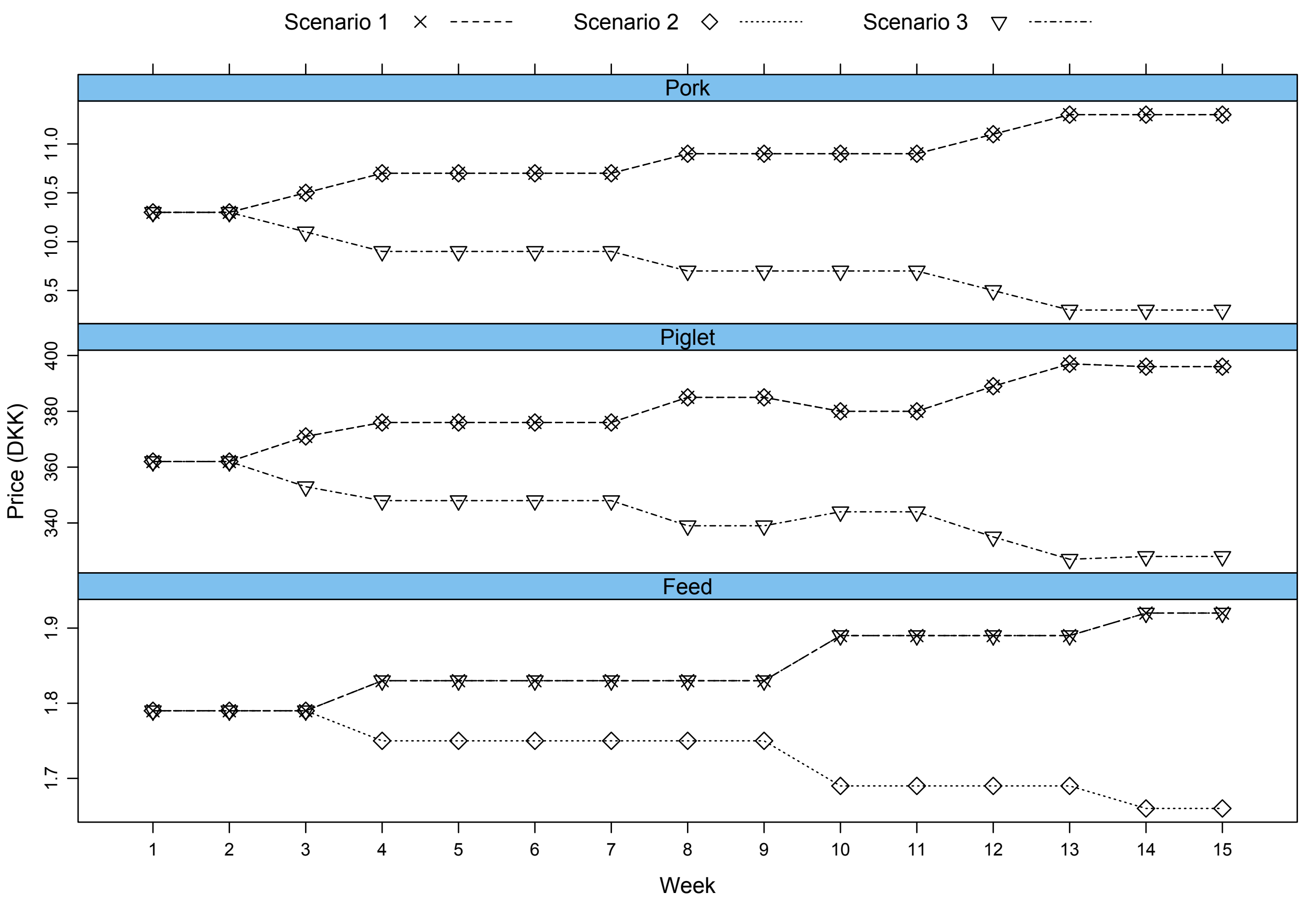

Figure 2: Price fluctuations in the three scenarios. In Scenario 1, the trends of feed and piglet prices are unfavourable and the trend of pork price is favourable. In Scenario 2, the trends of pork and feed prices are favourable and the trend of piglet price is unfavourable. In Scenario 3, the trends of pork and feed prices are unfavourable and the trend of piglet price is favourable.

Scenario 2: Favourable trends of pork and feed prices and unfavourable trend of piglet price. Pork price increases from 10.3 to 11.3 DKK, feed price decreases from 1.79 to 1.66 DKK and piglet price increases from 336 to 396 DKK.

Scenario 3: Unfavourable trends of pork and feed prices and favourable trend of piglet price. Pork price decreases from 10.3 to 9.3 DKK, feed price increases from 1.79 to 1.92 DKK and piglet price decreases from 362 to 328 DKK.

During the 15 weeks period, the average weight in the pen increases from 26.8 to 128.9 kg with a standard deviation increasing from 3 to 15.4 kg. Notice that the growth of the pigs is the same in the three scenarios and hence the only factor affecting the marketing policy is the price information.

The results for each scenario are illustrated in Figure 3. The optimal decision is shown just above the x-axis where the numbers denote the number of the heaviest pigs culled from the pen (decision "cull *q* pigs"), the letter "T" indicates the "terminate" decision , and the letter "C" corresponds to continuing the production process ("continue" decision). The bars show the number of remaining pigs in the pen before making a decision.

In Scenarios 1 and 2, fluctuations of pork and piglet prices are the same while the fluctuations in feed price are different. By comparing the two scenarios, we observe that the different trends of the feed price have a significant impact on the optimal policy. In Scenario 2, a decreasing feed price leads to an earlier termination (at week 11) compared to Scenario 1 with an increasing feed price (termination at week 15). Note that when the pen is terminated, a low feed price affects the feeding cost of the next production cycle and hence when the feed price is low, it may be beneficial to terminate the pen earlier and start a new production cycle. On the other hand, an increasing feed price (in Scenario 1) during the marketing period (an increase from 1.83 to 1.92 in weeks 9-15) results in a longer production cycle and individual marketings in weeks 11 to 14.

In Scenario 3, we have an increasing trend in feed prices (similar to Scenario 1) but unlike Scenarios 1 and 2, the trends of pork and piglet prices are decreasing in this scenario (see Figure 2). Here a decreasing piglet price does not result in an earlier termination as we had in Scenario 2. Like in Scenario 1, the termination occurs at week 15 in this scenario too, which is due to the increasing trend of the feed price. That is,

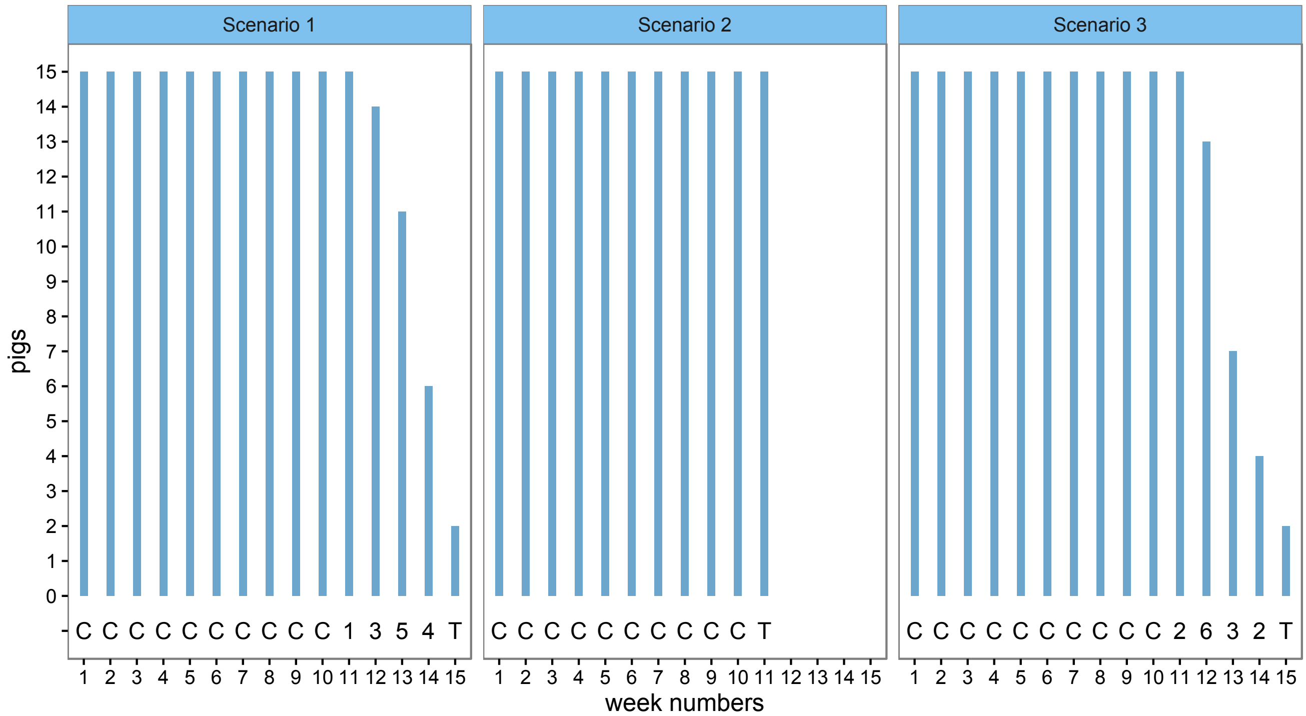

Figure 3: Optimal decisions of the HMDP for the three scenarios. The optimal decision is shown just above the x-axis where the numbers denote the number of the heaviest pigs culled from the pen (decision "cull q pigs"), the letter "T" indicates the "terminate" decision, and the letter "C" corresponds to continuing the production process ("continue" decision). The bars show the number of remaining pigs in the pen before making a decision.

the feed price compared to the piglet price has a higher impact on the optimal policy and reward. This observation was also supported in Pourmoayed, Nielsen, and Kristensen (2016). We also see that in Scenario 3 the fraction of remaining pigs in the pen in every week of the marketing period is lower than Scenario 1. This is because of the increasing trend of pork price in Scenario 1 that makes it more beneficial to keep more pigs in the pen and sell them in the next weeks while in Scenario 3 it is better to sell the pigs earlier since the pork price decreases in the next weeks.

Acknowledgement

This paper is funded by PigIT, a strategic research alliance founded in 2011 based on a grant from the Danish Council for Strategic Research (The PigIT project, Grant number 11-116191). The alliance has three partners (University of Copenhagen, Aarhus University and The Pig Research Centre).

References

• Danish Agriculture and Food Council. Statistics 2014 pigmeat. Technical report, 2015. URL http://www.agricultureandfood.dk/Prices_Statistics/Annual_Statistics.aspx.

• A.R. Kristensen and E. Jørgensen. Multi-level hierarchic markov processes as a framework for herd management support. Annals of Operations Research, 94(1-4):69–89, 2000.

• Landbrug & Fødevarer. Udenrigshandel 2010 - 2014 fødevareklyngen – varer til hele verden. Technical report, 2015. URL https://www.lf.dk/Tal_og_Analyser/Fodevareklyngens_udenrigshandel.aspx.

• L.R Nielsen. Mdp: Markov decision processes in R. R package v1.1., 2009. URL http://r-forge.r-project.org/projects/mdp/.

• L.R. Nielsen and A.R. Kristensen. Finding the K best policies in a finite-horizon Markov decision process. European Journal of Operational Research, 175(2):1164–1179, 2006.

• L.R. Nielsen and A.R. Kristensen. Markov decision processes to model livestock systems. In Lluis M. Plà-Aragonés, editor, Handbook of Operations Research in Agriculture and the Agri-Food Industry, volume 224 of International Series in Operations Research & Management Science, pages 419–454. Springer, 2014.

• R. Pourmoayed and L.R. Nielsen. An overview over pig production of fattening pigs with a focus on possible decisions in the production chain. Technical Report PigIT Report No. 4, Aarhus University, 2014. URL http://pigit.ku.dk/publications/PigIT-Report4.pdf.

• R. Pourmoayed and L.R. Nielsen. Github repository: Slaughter pig marketing under price fluctuations, 2015. URL https://github.com/pourmoayed/hmdpPricePigIT.git.

• R. Pourmoayed, L. R. Nielsen, and A. R. Kristensen. A hierarchical markov decision process modeling feeding and marketing decisions of growing pigs. European Journal of Operational Research, 250(3):925–938, 2016.

• Reza Pourmoayed. Optimization Methods in a Stochastic Production Environment. PhD thesis, Institut for Økonomi, Aarhus Universitet, 2016. URL http://pure.au.dk/portal/files/100477051/Reza_Pourmoayed_PhD_dissertation.pdf.

• R Core Team. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria, 2015. URL http://www.R-project.org/.

• H.C. Tijms. A first course in stochastic models. John Wiley & Sons Ltd, 2003. ISBN 978-0-471-49880-3.

• M. West and J. Harrison. Bayesian Forecasting and Dynamic Models (Springer Series in Statistics). Springer-Verlag, February 1997.

Reza Pourmoayed is a postdoctoral researcher at Operations Management group, Dept. of Engineering, Aarhus University. His research interests are stochastic optimization, operations research in agriculture, statistical modelling and machine learning.

Lars Relund Nielsen is Associate Professor at CORAL, Dept. of Economics and Business Economics, Aarhus BSS, Aarhus University. His interests include Markov decision processes, network algorithms and multi-objective optimization.