Textbook notes of herd management:

Basic concepts

Dina Notat No. 48 • August 1996

Anders R. Kristensen[1] and Erik Jørgensen[2]

This report is also available as a PostScript file on World Wide Web at URL: ftp://ftp.dina.kvl.dk/pub/Dina-reports/notatXX.ps

---

[1]Dina KVL
Dept. of Animal Science and Animal Health
Royal Veterinary and Agricultural University
Bülowsvej 13
DK-1870 Frederiksberg C
Tel: + 45 35 28 30 91
Fax: +45 35 28 30 87
E-mail: ark@dina.kvl.dk
WWW: http://www.dina.kvl.dk/~ark

[2]Dina Foulum
Department of Biometry and Informatics
Research Centre Foulum
P.O. Box 23
DK-8830 Tjele
Tel: + 45 89 99 12 30
Fax: +45 89 99 16 69
E-mail: ej@dina.sp.dk
WWW: http://www.sp.dk/~ej

Contents:

1. A general framework for herd management

Several points of view may be taken if we want to describe a livestock production unit. An animal nutritionist would focus on the individual animal and describe how feeds are transformed to meat, bones, tissues, skin, hair, embryos, milk, eggs, manure etc. A physiologist would further describe the roles of the various organs in this process and how the transformations are regulated by hormones. A biochemist would even describe the basic processes at molecular level.

A completely different point of view is taken if we look at the production unit from a global or national point of view. The individual production unit is regarded only as an arbitrary element of the whole livestock sector, which serves the purpose of supplying the population with food and clothing as well as manager of natural resources. A description of a production unit at this level would focus on its resource efficiency in food production and its sustainability from an environmental and animal welfare point of view.

Neither of these points of view are relevant to a herd management scientist even though several elements are the same. The herd management scientist also considers the transformation of feeds to meat, bones, tissues, skin, hair, embryos, milk, eggs and manure like the animal nutritionist, and he also regards the production as serving a purpose as we do at the global or national level. What differs, however, is the farmer. From the point of view of a herd management scientist, the farmer is in focus and the purpose of the production is to provide the farmer (and maybe his family) with as much welfare as possible. In this connection welfare is regarded as a very subjective concept and has to be defined in each individual case. The only relevant source to be used in the determination of the definition is the farmer himself.

The herd management scientist assumes that the farmer concurrently tries to organize the production in such a way that his welfare is maximized. In this process he has some options and he is subjected to some restraints. His options are to regulate the production in such a way that his welfare is maximized given the restraints. The way in which he may regulate production is by deciding what factors he wants to use at what levels. A factor is something which is used in the production, i.e. the input of the transformation process. In livestock production, typical factors include buildings, animals, feeds, labor and medicine. During the production process, these factors are transformed into products which in this context include meat, offspring, milk, eggs, fur, wool etc. The only way a farmer is able to regulate the production - and thereby try to maximize his welfare - is by adjustments of these factors.

Understanding the factors and the way they affect production (i.e. the products and their amount and quality) is therefore essential in herd management. Understanding the restraints is, however, just as important. What the restraints do is actually to limit the possible welfare of the farmer. If there were no restraints any level of welfare could be achieved. In a real world situation the farmer faces many kinds of restraints. There are legal restraints regulating aspects like use of hormones and medicine in production, storing and use of manure as well as housing in general. He may also be restricted by production quotas. An other kind of restraints are of economic nature. The farmer only has a limited amount of capital at his disposal, and usually he has no influence on the prices of factors and products. Furthermore,

he faces some physical restraints like the capacity of his farm buildings or the land belonging to his farm, and finally his own education and skills may restrict his potential welfare.

In general, restraints are not static in the long run: Legal regulations may be changed, quota systems may be abolished or changed, the farmer may increase or decrease his capital, extend his housing capacity (if he can afford it), buy more land or increase his mental capacity by training or education. In some cases (e.g. legal restraints) the changes are beyond the control of the farmer. In other cases (e.g. farm buildings or land) he may change the restraints in the long run, but has to accept them at the short run.

We are now ready to define herd management:

> *Herd management is a discipline serving the purpose of concurrently ensuring that the factors are combined in such a way that the welfare of the individual farmer is maximized subject to the restraints imposed on his production.*

The general welfare of the farmer depends on many aspects like monetary gain (profit), leisure time, animal welfare, environmental sustainability etc. We shall denote these aspects influencing the farmer's welfare as *attributes*. It is assumed that the consequences of each possible combination of factors may be expressed by a finite number of such attributes and that a uniquely determined level of welfare is associated with any complete set of values of these attributes. The level of welfare associated with a combination of factors is called the *utility value*. Thus the purpose of herd management is to maximize the utility value. A function returning the utility value of a given set of attributes is called a *utility function*. In Chapter 2, the concept of utility is discussed more thoroughly.

The most important factors in livestock production include:
- farm buildings
- animals
- feeds
- labor
- medicine and general veterinary services
- management information systems and decision support systems
- energy

In order to be able to combine these factors in an optimal way it is necessary to know their influence on production. As concerns this knowledge, the herd management scientist depends on results from other fields like agricultural engineering, animal breeding, nutrition and preventive veterinary medicine. The knowledge may typically be expressed by a production function, f, which in general for a given *stage* (time interval), *t*, takes the form:

$$\mathbf{Y}_{s,t} = \mathbf{f}_{s,i}(\mathbf{x}_{s,t}, \mathbf{x}_{s,t-1}, \ldots, \mathbf{x}_{s,1}) + \mathbf{e}_t \ , \qquad \textbf{(1)}$$

where $\mathbf{Y}_{s,t}$ is a vector of *n* products produced, $f_{s,i}$ is the production function, $\mathbf{x}_{s,t}$ is a vector of *m* factors used at stage *t* and $\mathbf{e}_t$ is a vector of *n* random terms. The function $f_{s,i}$ is valid for a given production unit *s*, which may be an animal, a group or pen, a section or the entire herd. The characteristics of the production unit may vary over time, but the set of observed characteristics at stage *t* are indicated by the *state* of the unit denoted as *i*. The state

specification contains all relevant information concerning the production unit in question. If the function is defined at animal level the state might for instance contain information on the age of the animal, the health status, the stage of reproductive cycle, the production level etc. In some cases, it is also relevant to include information on the disease and/or production history of the animal (for instance the milk yield of *previous* lactation) in the state definition.

The total production $\mathbf{Y}_t$ and factor consumption $\mathbf{x}_t$ is calculated simply as

$$\mathbf{Y}_t = \sum_{s \in S} \mathbf{Y}_{s,t} \tag{2}$$

and

$$\mathbf{x}_t = \sum_{s \in S} \mathbf{x}_{s,t} \ , \tag{3}$$

where S is the set off all production units s at the same level.

Eq. (1) illustrates that the production is only partly under the control of the manager, who decides the levels of the factors at various stages. The direct effects of the factors are expressed by the production function f, but the actual production also depends on a number of effects outside the control of the manager as for instance the weather conditions and a number of minor or major random events like infection by contagious diseases. These effects outside the control of the manager will appear as random variations which are expressed by $\mathbf{e}_t$. This is in agreement with the general experience in livestock production that even if exactly the same factor levels were used in two periods, the production would nevertheless differ between the periods.

An other important aspect illustrated by Eq. (1) is the *dynamic* nature of the herd management problem. The production at stage t not only depends on the factor levels at the present stage, but it may very well also depend on the factor levels at previous stages. In other words, the decisions made in the past will influence the present production. Obvious examples of such effects is the influence of feeding level on the production level of an individual animal. In dairy cattle, for instance, the milk yield of a cow is influenced by the feeding level during the rearing period, and in sows the litter size at weaning depends on the feeding level in the mating and gestation period.

The production $\mathbf{Y}_t$ and factors $\mathbf{x}_t$ used at a stage are assumed to influence the attributes describing the welfare of the farmer. We shall assume that k attributes are sufficient and necessary to describe the welfare. If we denote the values of these attributes at a specific stage t as $u_{1,t},...,u_{k,t}$, we may logically assume that they are determined by the products and factors of the stage. In other words, we have:

$$\mathbf{u}_t = \mathrm{h}(\mathbf{Y}_t, \mathbf{x}_t) \ , \tag{4}$$

where $\mathbf{u}_t = (u_{1,t},...,u_{k,t})$ is the vector of attributes. We shall denote h as the *attribute function*. The over-all utility of U_N for N stages, $t_1,...,t_N$, may in turn be defined as a function of these attributes:

$$U_N = \mathbf{g}(\mathbf{u}_{t_1},\mathbf{u}_{t_2},\ldots,\mathbf{u}_{t_N}) = \mathbf{g}(u_{1,t_1},\ldots,u_{1,t_N},u_{2,t_1},\ldots,u_{2,t_N},u_{k,t_1},\ldots,u_{k,t_N}) \ , \tag{5}$$

where g is the *utility function*, and N is the relevant time horizon. If we substitute Eq. (4) into Eq. (5), we arrive at:

$$U_N = \mathbf{g}(\mathbf{h}(\mathbf{Y}_{t_1},\mathbf{x}_{t_1}),\mathbf{h}(\mathbf{Y}_{t_2},\mathbf{x}_{t_2}),\ldots,\mathbf{h}(\mathbf{Y}_{t_N},\mathbf{x}_{t_N})) \ . \tag{6}$$

From Eq. (6) we observe, that if we know the attribute function h and the utility function g and, furthermore, the production and factor consumption at all stages have been recorded, we are able to calculate the utility relating to any time interval *in the past*. Recalling the definition of herd management, it is more relevant to focus on the *future* utility derived from the production. The only way in which the farmer is able to influence the utility is by making decisions concerning the factors. Having made these decisions, the factor levels $\mathbf{x}_t$ are known also for future stages. The production levels $\mathbf{Y}_t$, however, are unknown for future stages. The production function may provide us with the expected level, but because of the random effects represented by $\mathbf{e}_t$ of Eq. (1), the actual levels may very well deviate from the expected. An other source of random variation is the future state i of the production unit. This becomes clear if we substitute Eq. (1) into Eq. (5) (and for convenience assume that the production function is defined at herd level so that $s = S$):

$$U_N = \mathbf{g}(\mathbf{h}(\mathbf{f}_{s,i_1}(\mathbf{x}_{t_1},\ldots,\mathbf{x}_1)+\mathbf{e}_{t_1},\mathbf{x}_{t_1}),\mathbf{h}(\mathbf{f}_{s,i_2}(\mathbf{x}_{t_2},\ldots,\mathbf{x}_1)+\mathbf{e}_{t_2},\mathbf{x}_{t_2}),\ldots,\mathbf{h}(\mathbf{f}_{s,i_N}(\mathbf{x}_{t_N},\ldots,\mathbf{x}_1)+\mathbf{e}_{t_N},\mathbf{x}_{t_N})) \ . \tag{7}$$

From Eq. (7) we conclude, that even if all functions (production function, attribute function and utility function) are known, and decisions concerning factors have been made, we are not able to calculate numerical values of the utility relating to a future period. If, however, the distributions of the random elements are known, the *distribution* of the future utility may be identified.

When the farmer makes decisions concerning the future use of factors he therefore does it with incomplete knowledge. If, however, the distribution of the possible outcomes is known he may still be able to make rational decisions as discussed in the following chapters. Such a situation is referred to as *decision making under risk*.

2. Objectives of production: Farmer's preferences

A general characteristic of an attribute is that it directly influences the farmer's subjectively defined welfare and, therefore, is an element of the very purpose of production. This may be illustrated by a few examples. The average milk yield of the cows of a dairy herd is *not* an attribute of a utility function, because such a figure has no direct influence on the farmer's welfare. If, however, he is able to sell the milk under profitable conditions he will experience a monetary gain, which certainly may increase his welfare and, therefore, may be an attribute. In other words, the purpose of production from the farmer's point of view could never be to produce a certain amount of milk, but it could very well be to attain a certain level of monetary gain.

Animal welfare may in some cases be an attribute of the utility function. Whether or not it is in the individual case depends on the farmer's reasons for considering this aspect. An argument could be that animals at a high level of welfare probably also produce at a higher level and thereby increases the monetary gain. In that case, animal welfare is just considered as a short cut to higher income, but it is not considered to be a quality by itself. Accordingly, it should not be considered to be an attribute. If, on the other hand, the farmer wants to increase animal welfare *even if* it, to some extent, decreases the levels of other attributes like monetary gain or leisure time then it is certainly relevant to consider it to be an attribute of the utility function.

This discussion also illustrates that attributes are individual. It is not possible to define a set of attributes that apply to all livestock farmers. In the following section, however, we shall take a look at some examples of *typical* attributes of farmers' utility functions.

2.1. Common attributes of farmer's utility functions

Typical attributes describing the welfare of a livestock farmer include:
- Monetary gain
- Leisure time
- Animal welfare
- Working conditions
- Environmental preservation
- Personal prestige
- Product quality

In this section, we shall discuss each of these attributes and focus on how to define relevant attribute functions in each case and how to compare contributions from different stages.

2.1.1. Monetary gain

It seems very unlikely that monetary gain should not be an attribute of all professional farmers' utility function. In many practical cases it is even the only one considered when decisions are made. Usually, it is very easy to define the attribute function of monetary gain. If we assume product prices $p_{y1},...,p_{yn}$ and factor prices $p_{x1},...,p_{xm}$ to be fixed and known, the partial attribute function simply becomes (assuming that monetary gain is the 1st attribute):

$$u_{1t} = \mathrm{h}_1(\mathbf{Y}_t,\mathbf{x}_t) = \sum_{i=1}^{n} p_{yi}Y_{ti} - \sum_{i=1}^{m} p_{xi}x_{ti} \ . \qquad \textbf{(8)}$$

Monetary gains from different stages are not directly comparable. If a certain amount of money is gained at the present stage, it is possible to invest it and earn interest, so that the amount has increased at the following stage. The usual way to account for that is by discounting, so that we deal with the *present value* of future monetary gains. If we assume all stages to be of equal length and the interest rate to be constant, the present value of monetary gains from N stages is calculated as:

$$v_N = \sum_{t=1}^{N} \beta^{t-1} u_{1t} \ , \tag{9}$$

where $0 < \beta \leq 1$ is the *discount factor* usually calculated as

$$\beta = \exp(-r) \ , \tag{10}$$

where r is the interest rate pr. stage. For further information on the theory and principles of discounting, reference is made to textbooks on economics.

2.1.2. Leisure time

Just like monetary gain, leisure time is probably an attribute of all farmer's utility function. It is, however, slightly more complicated to represent it, because not only the total value, but also the distribution over time (day, week, year) is relevant. Since leisure time is the counterpart of work, it may therefore be relevant to split up the factor labor into several sub-factors like for instance work on weekdays from 6 a.m. to 6 p.m. x_1, work on weekdays from 6 p.m. to 6 a.m. x_2, work in weekends from 6 a.m. to 6 p.m. x_3, work in weekends from 6 p.m. to 6 a.m. x_4 etc. These sub-factors are assumed to refer to the farmer's own work, whereas the work by employees has to be represented by other sub-factors.

Using these sub-factors we may quite easily at any stage t calculate corresponding sub-attributes like total leisure time u_{2t1}, leisure time in weekends u_{2t2}, during day time u_{2t3} etc. If the stage length in days is denoted as D, these sub-attributes are objectively calculated as for instance total leisure time in hours:

$$u_{2t1} = 24D - \sum_j x_j \ , \tag{11}$$

and correspondingly for other sub-attributes. There exists no general over-all partial attribute function h_2 for leisure time, because the properties of that function depends on the farmer's individual (subjective) preferences. It seems however, reasonable to assume, that it may defined in each individual case as a function of a number, ν, of objectively calculated sub-attributes like those discussed, i.e.

$$u_{2t} = h_2(u_{2t1},\ldots,u_{2t\nu}) \ . \tag{12}$$

As concerns the comparison of values relating to different stages, the sub-attributes are to some extent additive (e.g. total leisure time). In other cases it may be more relevant to look at the average value (e.g. work load in weekends) and/or the maximum value. If, for instance, x_3 and x_4 are both zero, it means that the farmer does not have to work at all during the weekend, which some farmers would give a high priority. Correspondingly, if $u_{2t1} = 24D$ for a stage, it means that the farmer is able to have a vacation.

2.1.3. Animal welfare

While monetary gain and leisure time are probably attributes of all farmers' utility functions

(albeit that the relative weights may differ), the situation is different with animal welfare. There is no doubt that some farmers will include the attribute, but since the level of animal welfare has no direct consequences for the farmer's welfare it is probably ignored by others. It does not mean that the farmer necessarily ignores animal welfare, but the reason for considering it may only be for legal reasons or because it to some extend affects the production and consequently other attributes like monetary gain and/or labor. If that is the case, animal welfare is only considered as far as to meet legal demands and not to affect the other attributes. In such a situation it should not be defined as a separate attribute of the utility function.

Recalling Section 2.1.2. we note that monetary gain is an attribute which may objectively be represented by a single numerical value. As concerns leisure time, we needed several numerical values (or sub-attributes) for an objective representation. If we want to express that attribute as a single numerical value it is unavoidable a subjective personal value derived from the preferences of that particular farmer. An other farmer would probably attach different weights to the various objective sub-attributes involved and, therefore, his over-all evaluation of the scenario would differ.

When it comes to an attribute like animal welfare, it is even a question whether objective sub-attributes exist. In general, the assessment of the very concept of animal welfare is a problem. In the literature many different definitions have been given. A few examples collected by Kilgour & Dalton (1984) are listed for illustration:

- *A wide term that embraces both the physical and mental wellbeing of the animal.* (Brambell, 1965).
- *Existence in reasonable harmony with the environment, both from an ethological and physiological point of view.* (Dutch National Council for Agricultural Research, 1977).
- *Welfare can be satisfied if three questions can be answered in the affirmative: (a) Are the animals producing normally? (b) Are they healthy and free from injury? (c) Is the animal's behaviour normal?* (Adler, 1978).
- *Welfare is a relative concept. Profit is a matter related to welfare and determining the* ***relationship*** *between welfare and profit is a scientific matter. The* ***choice*** *between welfare and profit is an ethical matter.* (Brantas, 1975).
- Considering that while an animal is producing protein without observable signs of pain, then it can be considered to be comfortable. *Distress, strain, gross abuse and suffering are used to describe unfavouable circumstances ... the presence of raw flesh or heavy bruises would be classified this way, but callouses or hard skin caused by concrete floors and producing no pain would not be classed as distressing.* (Randall, 1976).
- *On a general level it is a state of complete mental and physical health where the animal is in harmony with its environment. On an empirical level it may be measured by studying an animal in an environment which is assumed to be ideal, and then comparing it with an animal in the environment under investigation.* (Hughes, 1976).
- *Handling animals in the least disturbing manner with full consideration for their normal species-specific behaviour requirements.* (Kilgour, 1987).

None of the definitions listed are appropriate in this context. A more consistent approach seems to be to consider animal welfare analogously to human welfare as discussed by Sandøe & Simonsen (1992). If we accept this approach, the welfare of an animal may be described

by a utility function involving a number of attributes each representing an element of animal welfare. The principles are exactly the same as with the farmer's utility function, but the attributes considered naturally differ. Animal welfare has nothing to do with monetary gain, environmental preservation or product quality. Instead, relevant attributes might be:

- Absence of hunger
- Absence of pain
- Thermal comfort
- Absence of fear
- Pleasure, joy

Just like the farmer's utility, animal welfare may be controlled by the dynamic allocation of factors to the production processes. Obviously, absence of hunger is ensured by an appropriate supply of feeds, but it is important to realize that what matters in relation to welfare is the subjective experience by the animal. It is certainly possible to feed animals in such a way, that all requirements for maintenance and production are satisfied, but the animal nevertheless feels hungry. Such a situation may occur if an animal is fed a very concentrated ration as it is typically the case with sows during the gestation period. Even though all physiological concerns are met, the welfare of the animal may be violated.

As concerns pain, many factors are involved: Physical injuries caused by inapropriate housing conditions may expose the animal to pain. The same applies of course to injuries caused by other animals in the flock. An other source of pain is disease, which may be prevented or cured by drugs, vitamins, minerals or feeding in general as well as by the housing conditions.

By thermal comfort we mean that the animal neither feels cold or too warm. This may be achieved in many different ways involving many different factors. If an animal feels cold, we may for instance install a heating system, reduce the ventilation, increase the number of animals per square meter, or supply the animal with straw for bedding.

It seems reasonably to assume that farm animals are able to experience fear. If we accept this assumption we have to consider how to prevent it. The two most important factors in this respect are probably labor and other animals. A negative influence from labor is prevented by appropriate daily routines and the influence from other animals is regulated through flock sizes and re-grouping.

As concerns the last attribute mentioned (pleasure, joy) it is a question whether it makes sense in relation to animals, but according to Sandøe & Simonsen (1992) several researchers define this attribute and even attributes like satisfaction and expectation as elements of animal welfare.

Having identified the attributes is of course only half the way to an operational representation of animal welfare as a possible attribute of the farmer's utility function in relation to herd management. Since the attributes of animal welfare represent the animals' subjective experience, they cannot be measured directly. All that can be measured are physiological, behavioral, pathological and other objective parameters which may serve as evidence for the occurrence of the relevant subjective experiences. Sandøe & Simonsen (1992) therefore argue that animal welfare researchers have to find out which measurable parameters will serve as

indicators of the occurrence of which experiences. This is an other example of a situation where the herd management scientist depends on results achieved in other research disciplines. What is actually measured are typically parameters like hormone levels, conflict and abnormal behavior, results of various sorts of choice tests conducted and disease incidences. Sandøe & Simonsen (1992) refer to Dawkins (1980) and Fraser & Broom (1990) for a more full and detailed account.

Reference is also made to Sandøe & Hurnik (1996) for a collection of articles discussing concepts, theories and methods of measurement in relation to the welfare of domestic animals.

An other problem is how to combine the conclusions regarding individual attributes into the aggregate notion of over-all animal welfare, which is the possible attribute of the farmer's utility function. Since, however, this problem is analogous to the problem of combining the various attributes of the farmer's utility function, reference is made to Section 2.2., where that issue is discussed.

2.1.4. Working conditions

Also the working conditions may affect the farmer's welfare. When working in barns, dust may contaminate the respiratory system and even damage the lungs. Correspondingly, very hard work (heavy burdens etc.) may harm the back. It is probably not possible to express the working conditions objectively by a single numerical attribute. It must, however, be assumed that if the general working conditions are so bad that they are a threat to the farmers permanent health condition improvements will have a very high priority in the utility function. When expressing the over-all value of the attribute over several stages it may therefore be relevant to define it as (assuming that the working conditions are the 4*th* attribute of the utility function):

$$u_{4N} = \min_{t}\{u_{4t}\} \ , \qquad \textbf{(13)}$$

where u_{4t}, $1 \le t \le N$, is the resulting value at stage t of the attribute representing the working conditions. In words, Eq. (13) says that what matters to the farmer are the *worst* working conditions experienced during the N stages.

2.1.5. Environmental preservation

Just like some of the previously discussed attributes it may also be relevant initially to define a number of sub-attributes $u_{5t1},...,u_{5t}$ each representing an aspect of environmental preservation. In most cases they may be objectively calculated as numerical figures even though the precision may vary. An example of such a sub-attribute is according to Jensen & Sørensen (1996) the net loss of nutrients (e.g. nitrogen) to the environment which may be calculated as

$$u_{5t1} = \sum_{i=1}^{m} c_{xi} x_{ti} - \sum_{i=1}^{n} c_{Yi} Y_{ti} \ , \qquad \textbf{(14)}$$

where c_{xi} and c_{Yi} are the concentrations of the nutrient in question in the ith factor and the ith product, respectively. Other relevant sub-attributes mentioned by Jensen & Sørensen (1996)

include the total consumption of energy and indicators of the risk of residues of pesticides in the environment.

When dealing with the leisure time attribute we concluded that no general over-all partial attribute function exists. For an individual farmer, however, we assumed that such a subjectively defined function exists as illustrated by Eq. (12). As concerns environmental preservation, the aggregate attribute function probably also has to be subjectively defined, but, in this case the reason is merely lack of knowledge on the relative importance of the sub-attributes in relation to environmental preservation.

When the attribute is evaluated over several stages, calculation of average values of the sub-attributes are most often the relevant method.

2.1.6. Personal prestige

Already by using the word "personal" we suggest that we are dealing with a subjectively defined attribute. In many cases, personal prestige has something to do with a high level of productivity (e.g. milk yield per dairy cow or litter size per sow), the presence of a factor of a particular kind (e.g. a milking parlor or a big tractor) implying that we are dealing with sub-attribute functions of the kind

$$u_{6t1} = \frac{Y_{ti}}{x_{tj}}, \tag{15}$$

(where Y_{ti} and x_{tj} are, for instance, the total milk yield and the number of cows at stage t) or the kind

$$u_{6t2} = \begin{cases} 1, & x_{ti}>0 \\ 0, & x_{ti}=0 \end{cases} . \tag{16}$$

Even though the list of sub-attributes is highly subjective, the calculation of their numerical values are, as illustrated, most often objective. Otherwise their worth as the basis of personal prestige would be lacking. Also the over-all partial attribute function h_6 is subjectively defined as illustrated by Eq. (12).

2.1.7. Product quality

All farmers are probably interested in product quality to the extent that it influences the price of the product (e.g. the hygienic status of milk). If, however, that is the only reason for considering product quality, it should *not* be defined as an attribute of the farmer's utility function. If the farmer is interested in product quality also if it is independent of (or even in conflict with) monetary gain and other attributes already included then it is relevant to consider it as an attribute of the utility function.

The relevant sub-attributes to consider will of course depend on the product, but typical examples when dealing with animal production are the hygienic status, the risk of residues of medicine or indicators of the nutritional value for humans (e.g. fat or protein percentage).

2.2. From attributes to utility

Having identified the attributes of a farmer's utility function we face the problem of how to combine these single attributes into an aggregate over-all utility of the farmer. The examples of the previos section clearly illustrate that typical attributes are very different in nature and measured in different units (e.g. monetary gain and animal welfare). When we use the attributes in a planning situation we furthermore face the problem that we are not able to calculate exact values relating to future stages as discussed in relation to Eq. (7). We have to consider the attributes (and thereby utilities) of future stages as random variables reflecting that we are producing under risk. In other words, the farmer has to make decisions based on *distributions* of attributes and utilities rather than fixed values.

2.2.1. Single attribute situation: Risk

In the most simple case the farmer's utility function only include one attribute, i.e.

$$U_N = g(u_N) \ , \tag{17}$$

where u_N is the aggregate value of the attribute over N stages. For convenience, we shall assume that the attribute is monetary gain, but in principle it could be any attribute. In a real situation it is hardly likely that a utility function only depends on one attribute, but the considerations below are also relevant if we consider a decision that only influences one attribute, so that the remaining attributes may be considered as fixed.

In order to illustrate the nature of a typical single-attribute utility function we shall initially consider a question: *How do we feel about earning an extra fixed amount of money?* We shall denote our present annual income as a, and the potential extra income as Δa (where $\Delta a > 0$). There is probably no doubt that we would all prefer the total income $a + \Delta a$ to a no matter the actual values of a and Δa. In other words, the utility function g must have the property that

$$g(a+\Delta a) \geq g(a) \tag{18}$$

for $\Delta a > 0$.

In order to be able to answer the question more specificly, it is not sufficient information just to know that Δa is positive. We also have to know the value of not only Δa, **but also of *a***. A simple example illustrates this. If our present annual income is a million (i.e. a = 1,000,000) we are probably rather indifferent to earning an extra amount of say ten thousand (i.e. Δa = 10,000). If, on the other hand, our present annual income is only fifty thousand (a = 50,000), we would regard the same option as a considerable improvement of our situation. In other words, the marginal utility of the same monetary gain depends heavily on the present level of income. This is illustrated in Figure 1 which shows the typical course of a single attribute utility function defined as

$$U = u^{\alpha} \ , \tag{19}$$

where most often $0 < \alpha < 1$. In Figure 1, $\alpha = 1/3$. For a discussion of other common algebraic

representations of utility functions, reference is made to Anderson et al. (1977).

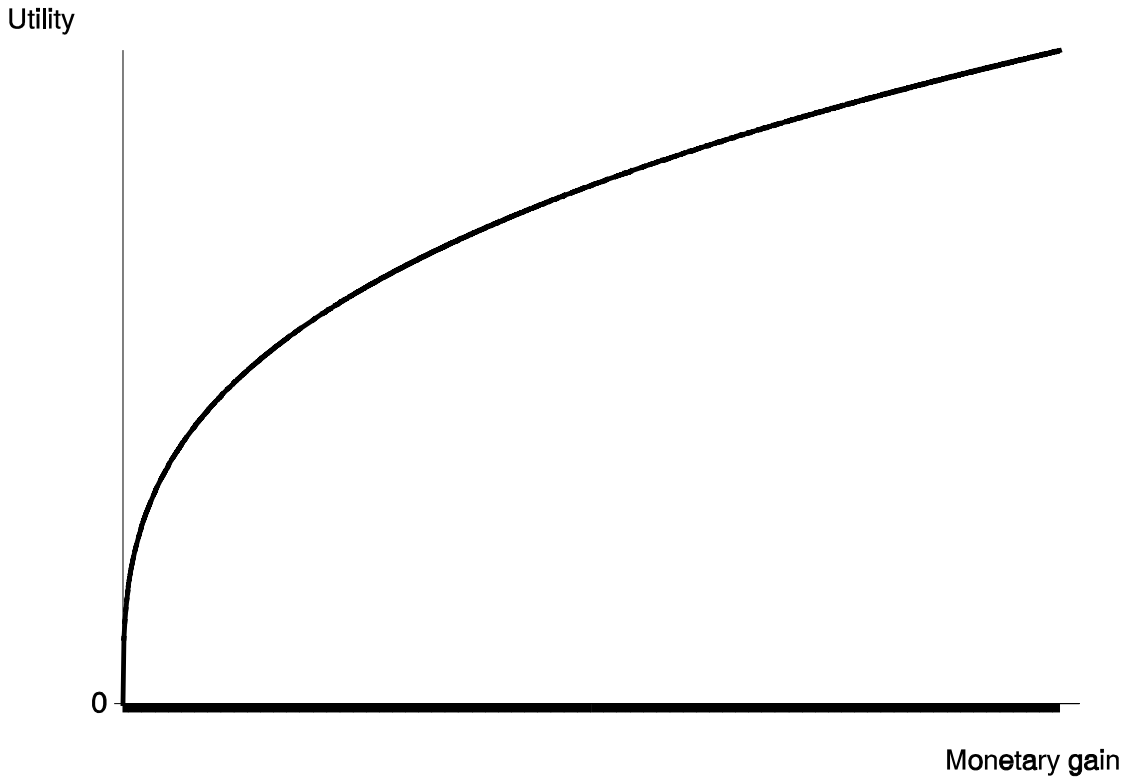

Figure 1. A typical single attribute utility function.

Now, assume that we have to choose between two alternative actions. Since actions always relate to the allotment of factors, we shall denote them as $\mathbf{x}^1$ and $\mathbf{x}^2$, respectively. The resulting production from these actions are correspondingly denoted as $\mathbf{Y}^1$ and $\mathbf{Y}^2$. Applying Eqs. (1) and (8) we may calculate the monetary gain under action 1 as:

$$u_1^1 = \sum_{i=1}^{n} p_{yi}\left(\mathrm{f}_i(\mathbf{x}^1) + e_i^1\right) - \sum_{i=1}^{m} p_{xi}x_i^1 = \sum_{i=1}^{n} p_{yi}\mathrm{f}_i(\mathbf{x}^1) - \sum_{i=1}^{m} p_{xi}x_i^1 + \varepsilon^1 , \tag{20}$$

where the random variable ε^1 is defined as

$$\varepsilon^1 = \sum_{i=1}^{n} p_{yi}e_i^1 . \tag{21}$$

The monetary gain under action 2 is calculated correspondingly. Assuming fixed prices, the only random elements of the monetary gains are ε^1 and ε^2. Without loss of generality, we may assume that $E(\varepsilon^1) = E(\varepsilon^2) = 0$. Eq. (20) may therefore be written as

$$u_1^1 = E(u_1^1) + \epsilon^1 . \tag{22}$$

Let us assume that $E(u_1^1) = 100$ and $E(u_1^1) = 105$. If the outcome is known with certainty (i.e. $P(\varepsilon^1=0) = P(\varepsilon^2=0) = 0$), the choice between the two actions is easy. Using the utility function defined in Eq. (19) we may calculate the corresponding utilities as $U^1 = 100^{1/3} = 4.64$ and $U^2 = 105^{1/3} = 4.72$ implying that action 2 is preferred over action 1. This is certainly not surprising, and in fact we only needed the logically deduced Eq. (18) in order to arrive at the same conclusion.

If, however, risk is involved the situation is more complicated. By risk we mean that the random variables ε^1 and ε^2 have known distributions with variances greater than zero. A simple numerical example shall illustrate how this may influence the expected utility associated with each of the two alternative actions. In Table 1, the assumed distributions of ε^1 and ε^2 are specified, and the resulting expected utilities are calculated. It should be emphasized that the expected monetary gains under the two actions are still 100 and 105, respectively.

Table 1. Probability distributions and expected utilities of two alternative actions.

Action 1				Action 2			
ε^1	u_1^1	$P(\varepsilon^1)$	U^1	ε^2	u_1^2	$P(\varepsilon^2)$	U^2
−10	90	0.25	4.48	-65	40	0.25	3.42
0	100	0.50	4.64	0	105	0.50	4.72
10	110	0.25	4.79	65	170	0.25	5.54
Expected utility, $E(U^1)$			4.64	Expected utility, $E(U^2)$			4.60

It may surprise to observe that even though the expected monetary gain under action 2 is higher than under action 1, the opposite applies to the expected utilities. Thus in the example we should indeed prefer action 1 over action 2. The explanation is simply that action 2 is more risky than action 1, and since the applied utility function only modestly rewards outcomes greater than the expected value but heavily punishes lower outcomes (cf. Figure 1), the expected utility of a risky action will always be lower than the expected utility of a less risky action having the same expected outcome.

The fact that maximization of expected utility is the relevant criterion for choosing between risky actions follows from the so-called *expected utility theorem* which is discussed by Anderson et al. (1977, p 65-69). The theorem may be deduced from only three perfectly reasonable axioms describing consistent human attitudes to risky choices. The same axioms also ensure the very existence of a unique subjective utility function.

A common numerical representation of the risk associated with an action is the variance of the outcomes. In the example we may directly calculate the variances from Table 1. Under action 1 the variance is $V^1 = 0.25 \times (-10)^2 + 0.50 \times 0^2 + 0.25 \times 10^2 = 50$. The corresponding variance under action 2 is $V^2 = 0.25 \times (-65)^2 + 0.50 \times 0^2 + 0.25 \times 65^2 = 2112.5$ indicating a far more risky action.

In the single attribute situation, it is only because of random variation (i.e. risk) that we have to consider the concept of utility and the shape of the function converting outcomes to utilities. In many cases it may therefore be more natural to use the expected monetary gain and the variance directly in the evaluation of risky actions. That approach is called (E,V) analysis. The idea is to draw so-called iso-utility curves in a diagram where the X axis is the variance and the Y axis is the expected monetary gain. An iso-utiliti curve has the property, that all loci along the curve represent mean-variance combinations that yield the same level of utility. Any risky action with known expected gain and variance may be plotted as a locus in the diagram. By comparing the loci of the actions with the course of the iso-utility curves, the action representing the highest expected utility may be chosen.

In Figure 2, iso-utility curves representing 3 constant levels of utility are shown in an (E,V) diagram. If we plot the two actions from the numerical example of Table 1 in the diagram (recalling that Action 1 had the lowest variance and the lowest expected monetary gain) we easily see that action 1 represents the highest expected utility.

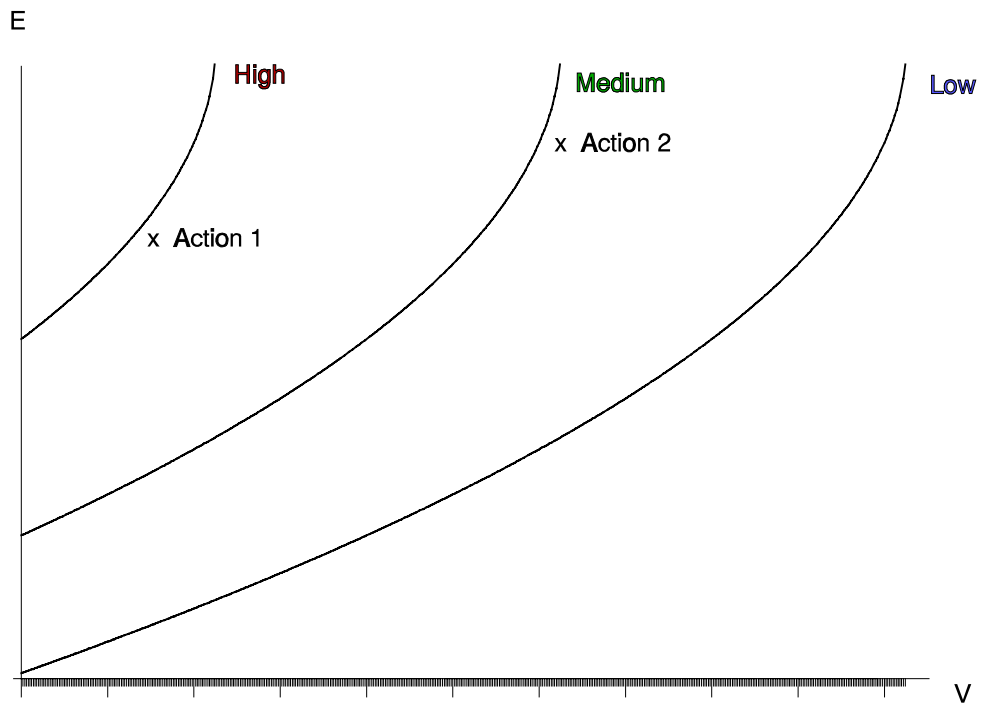

Figure 2. Iso-utility curves for 3 levels of utility. The two actions defined in Table 1 have been plotted in the diagram.

In the discussion of this section, we have implicitely assumed that the farmer in general

regards risk as something that should be minimized. In other words we assume him to be *risk averter*. The concave utility curve of Figure 1 represents a risk averter. At least in theory, a farmer may be *risk preferrer*. In that case his utility function is convex. Also in that case, Eq. (19) may be used as an algebraic representation of the utility function. For a risk preferrer, however, the parameter α has to be greater than 1. For a discussion of those aspects, reference is made to Anderson et al. (1977), whom we also refer to for a discussion of how to identify a farmer's utility function in practise.

2.2.2. Multiple attribute situation

In case that an action influences more than one attribute of the farmer's utility function we have to consider not only how each action affects each attribute, but also how much weight relatively the fulfilment of each attribute should be given. Animal breeding scientists face a completely analogous problem when they define breeding goals involving several traits as illustrated by a small example. Let us assume that the three traits *milk yield*, *daily gain* and *fertility* are relevant when dairy sires are selected for breeding. The relative breeding values for each trait may be estimated by usual methods, and we shall assume that we have to select one sire among the three candidates listed in Table 2.

Table **2**. Three sires and their relative breeding values for the traits *milk yield*, *daily gain* and *fertility*. Index 1 and Index 2 have been calculated using different economic weights.

Sire number	Milk yield	Daily gain	Fertility	Index 1	Index 2
1	110	107	104	107.75	107.00
2	109	106	103	106.75	106.00
3	114	103	102	108.25	105.50

Initially, we shal only consider the first four columns of Table 2. If we only had to choose between Sire 1 and Sire 2 there would be no problem since Sire 1 is superior in all three traits, but when we have to choose between Sire 1 and Sire 3 we face a problem, because Sire 3 is superior with respect to milk yield but inferior with respect to daily gain and fertility. There is no general solution to this problem, because the final choice will depend on the weight attached to milk yield relative to the weights attached to daily gain and fertility.

The animal breeding scientist solves the problem by defining a so-called breeding index which is a linear combination of the individually estimated breeding values, i.e.

$$B_i = \sum_{j=1}^{J} c_j b_{ij} \qquad \textbf{(23)}$$

where B_i is the breeding index of the *i*th animal, c_j is the weight attached to the *j*th trait, and b_{ij} is the estimated breeding value of the *i*th animal concerning the *j*th trait. If, for instance, we define the relative weight attached to milk yield to be 50% and the relative weights attached to daily gain and fertility to be 25% each (i.e. c_1 = 0.50 and c_2 = c_3 = 0.25) the

breeding indexes of the three sires of the example become as shown under *Index 1* in Table 2. As it appears from the table, Sire 3 should be chosen. If on the other hand, the relative weight attached to daily gain was 50% leaving 25% for each of the two remaining traits (i.e. $c_1 = c_3 = 0.25$ and $c_2 = 0.50$) the breeding indexes become as shown under *Index 2* leaving Sire 1 as the superior one. This clearly illustrates that the relative importance attached to the traits determines the mutual ranking of the animals.

In herd management science we do not choose among Sires but among alternative actions representing different factor allotments. Neither do we evaluate alternatives on breeding values but on attributes of the farmer's utility function. Nevertheless, if in the example we replaced *Sire* by *Action*, *trait* by *attribute* (*milk yield* by *monetary gain*, *daily gain* by *leisure time*, *fertility* by *animal welfare*) and *index* by *utility*, the example would still make sense. Also in herd management it is very likely that if we, for instance, rank actions on monetary gain the ranking changes if we use leisure time as our criterion of evaluation instead. Also the multi-attribute utility function may be defined analogously with Eq. (23). Using the notation of the previous sections, the utility function would just be a linear combination of the individual attribute functions:

$$U = \sum_{j=1}^{J} c_j u_j = \sum_{j=1}^{J} c_j \mathbf{h}_j(\mathbf{Y},\mathbf{x}) \ , \tag{24}$$

where $c_1,...,c_J$ are constant values which we may assume sum to 1 without loss of generality.

In the animal breeding example, Eq. (23) says that a higher breeding value concerning one trait may compensate for a lower value concerning other traits. The same is expressed for attributes like monetary gain, leisure time and animal welfare in Eq. (24). This is certainly a realistic property of a utility function, but, nevertheless, if the utility function is defined as in Eq. (24) we implicitly make some assumptions which by a closer look may seem unrealistic. In order to illustrate this we shall assume that only two attributes (monetary gain u_1 and leisure time u_2) are relevant for the management problem considered. Eq. (24) thus reduces to

$$U = cu_1 + (1-c)u_2 \ . \tag{25}$$

If we set U equal to some constant U', Eq. (25) may be rearranged into

$$u_2 = \frac{U'}{1-c} - \frac{c}{1-c}u_1 \tag{26}$$

implying that all combinations of u_1 and u_2 yielding the same utility U' form a linear relationship in a diagram. In Figure 3, three such *iso-utility curves* representing a low, medium and high fixed level of utility U' are shown. The course of the curves illustrate that decreased leisure time may be compensated by increased monetary gain, and the linear relationship implies that the marginal rate of substitution between the two attributes is constant. In other words, a constant improvement Δu_1 of the monetary gain may compensate a constant reduction Δu_2 of leisure time *no matter whether the initial level of monetary gain or leisure time is high or low*. From Eq. (26) we see that the constant marginal rate of substitution is:

$$\frac{\Delta u_1}{\Delta u_2} = \frac{c}{1-c} \; . \tag{27}$$

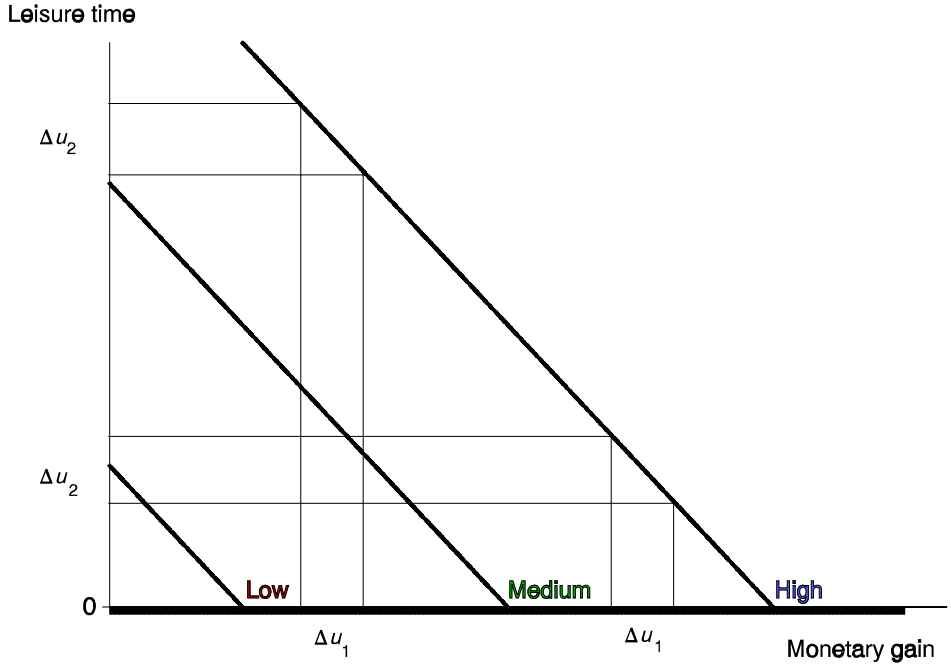

Figure 3. Iso-utility curves representing a low, medium and high fixed level of utility calculated according to Eq. (25). $\Delta u_2/\Delta u_1$ is constant.

If we consider the situation of a real farmer, it is hardly likely that his marginal rate of substitution is constant. If, in his present situtation, he has a high level concerning leisure time and a low level concerning monetary gain, an improvement concerning monetary gain will probably be given very high priority. The farmer is probably willing to sacrify quite a bit of his leisure time in order to improve his income. In the opposite situation (high level concerning monetary gain and low level concerning leisure time) he is probably only willing to sacrify very little leisure time in order to improve his monetary gain.

In Figure 4, iso-utility curves illustrating varying marginal rates of substitution between attributes are shown. Again, three fixed levels of utility are represented in the diagram. As shown in the figure, the marginal rate of substitution $\Delta u_{21}/\Delta u_1$ is very high if we are dealing with a poor farmer with plenty of leisure time whereas it is very low ($\Delta u_{22}/\Delta u_1$) in the opposite situation (a rich farmer working all the time).

A possible algebraic representation of a utility function having iso-utility curves like those of Figure 4 is:

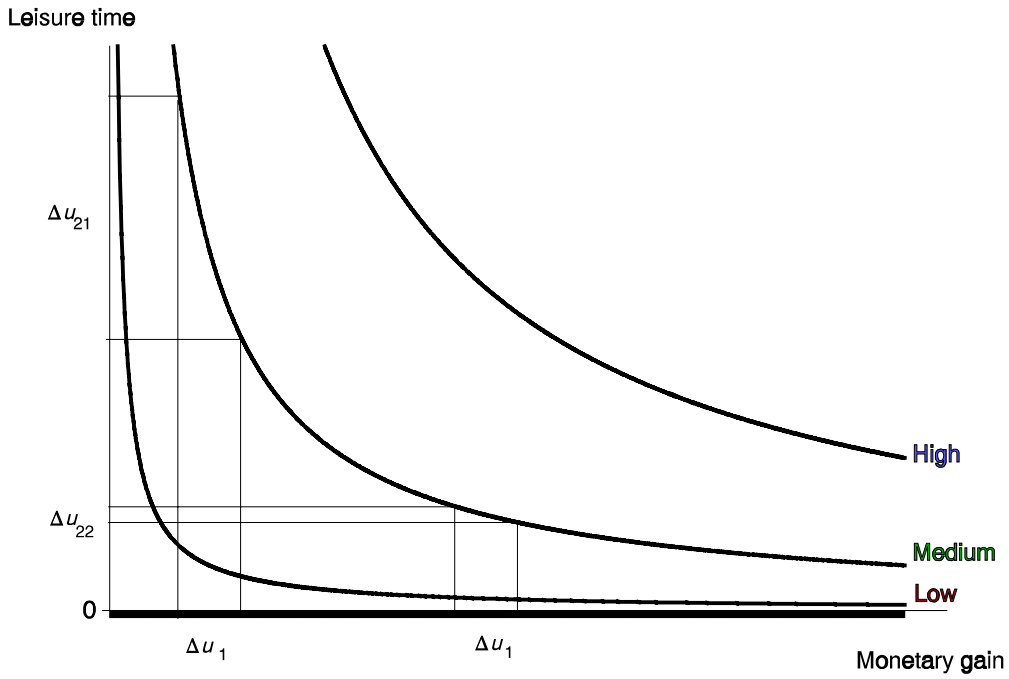

Figure 4. Iso-utility curves corresponding to the utility function of Eq. (28). $\Delta u_{21}/\Delta u_1 \gg \Delta u_{22}/\Delta u_1$.

$$U = u_1^{\alpha} u_2^{\beta} \ , \tag{28}$$

where (for the risk averter) $0 < \alpha < 1$ and $0 < \beta < 1$. It should be noticed, that if one of the attributes is left constant, Eq. (28) reduces to the single-attribute utility function of Eq. (19). For a discussion of other possible representations of a multi-attribute utility function, reference is made to Anderson et al. (1977).

2.2.3. Operational representation of a farmer's utility function

Tell me, how shall I feed my cows this winter? the farmer asks. The herd management scientist answers him back: *Let me know your utility function, and I shall tell you how to feed your cows.* From a theoretical point of view, the previous sections have illustrated that the herd management scientist is right. In order to choose an optimal action among a set of alternatives, we need to know the farmer's utility function. In other words we have to know what attributes he includes and how they are weighted in a specific situation. On the other hand, the farmer is off course not able to specify his utility function, so what do we do? Does the dialogue end here, or are we able to identify the utility function to an extent where it makes sense to make decisions?

Even though for instance Anderson et al. (1977) to some extend describe how single- and multi-attribute utility functions are identified in practice, it remains still a very difficult task and even the representation of relevant attribute functions for attributes like animal welfare

and working conditions may be a problem.

On the other hand, we may argue that it is not always necessary to know everything about the utility function when an action is chosen. In several cases we may logically conclude that incomplete information is sufficient to choose an action:

- In some cases, an action only influences one attribute. If risk is involved, it is sufficient to know the single attribute utility function. In case of a decision not involving risk, even the attribute function is sufficient.
- Many actions (typically those with short time horizon) only marginally influence the attributes. It is therefore sufficient to know the *local* marginal rates of substitution between attributes. In Figure 4, this may be illustrated by plotting the *current* combination of monetary gain and leisure time and draw a straight line with a slope equal to the marginal rate of substitution through that locus. It is not claimed that this constant rate is universal, but for marginal changes the approximation may suffice. This means that the simple additive Eq. (24) is used as a local approximation to a theoretically correct (but unknown) utility function with varying substitution rates.
- The utility function may only vary little over a rather large range of actions. It may be sufficient to choose a satisfactory action rather than an optimal one. This is particularly the case if a farmer's single-attribute utility function has a course as illustrated in Figure 5. In that case it is important, that the value of the attribute in question is at least as high as some satisfactory level u^*. Lower levels are punished by the utility function, but higher values are not rewarded. The attribute in question in Figure 5 could for instance be animal welfare. Most farmers will probably try to improve it if it is very low, but having reached a certain level, they do not worry about it anymore. When they make decisions with long time horizon (for instance regarding housing facilities) they make some decisions to ensure a satisfactory level, and afterwards they consider the matter of animal welfare to be out of concern.

An alternative to utility functions is the lexicographic utility concept which may be more operational in a real world situation. Instead of defining an aggregate multi-attribute utility function the farmer has to rank attributes from the most important to the least important. If four attributes are relevant the ranking could for instance be as follows:

1. Monetary gain (u_1)
2. Working conditions (u_2)
3. Animal welfare (u_3)
4. Leisure time (u_4)

In the pure form, the concept requires that actions are first evaluated entirely on the attributed given highest priority (in this case monetary gain), and action 1 is preferred to action 2 if and only if $u_1^1 > u_1^2$ (no matter the values of the other attributes). Only if $u_1^1 = u_1^2$, the actions are evaluated on the second attribute (working conditions) and, again, action 1 is preferred to action 2 if, and only if, $u_2^1 > u_2^2$ and so on. Only if $u_1^1 = u_1^2$, $u_2^1 = u_2^2$ and $u_3^1 = u_3^2$ the actions are evaluated on the fourth attribute (leisure time).

The lexicographic concept may also be combined with definition of satisfactory levels concerning the most important attributes and then maximizing the least important one subject to restraints on the others.

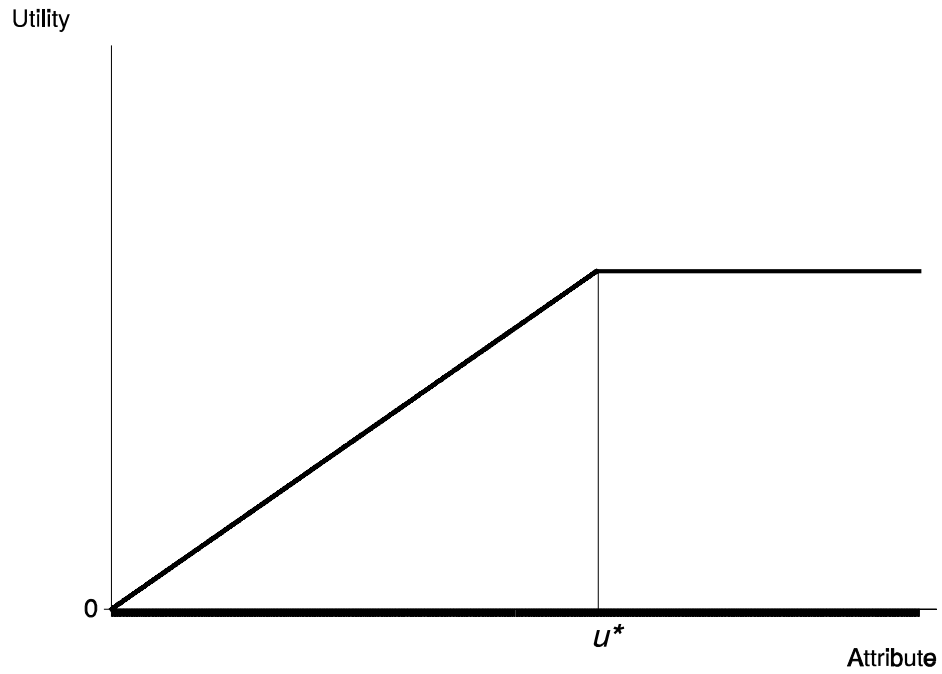

Figure 5. A single-attribute utility function reflecting the existence of a satisfactory level.

Finally, the obvious direct choice option should be mentioned. The farmer is just confronted with the multi-attribute consequences of the alternative actions. In other words, he is told that action 1 represents a monetary gain of u_1^1, working conditions expressed as u_2^1, animal welfare at the level u_3^1 and leisure time at the level u_4^1. Similar values concerning the other possible actions are given, and based on this information the farmer chooses the preferred action directly. This method has the obvious advantage that the advisor does not have to know the farmer's utility function, but only the relevant attributes.

3. The management cycle

3.1. The elements of the cycle

Herd management is a cyclic process as illustrated by Figure 6. The cycle is initiated by identification of the farmer's utility function as discussed in Chapter 2. Also the restraints have to be identified no matter whether they are of the legal, economic, physical or personal kind discussed in Chapter 1. The number of restraints will depend heavily on the time horizon considered. If the time horizon is short, the farmer faces more restraints of economic, physical and personal nature than when he is considering a long time horizon.

Having identified the farmer's utility function and the relevant restraints, one or more *goals* may be defined. It is very important not to confuse goals with the attributes of the utility function. The attributes represent basic preferences of the farmer, and they are in principle invariant, or - to be more precise - they only vary if the farmer's preferences change (for

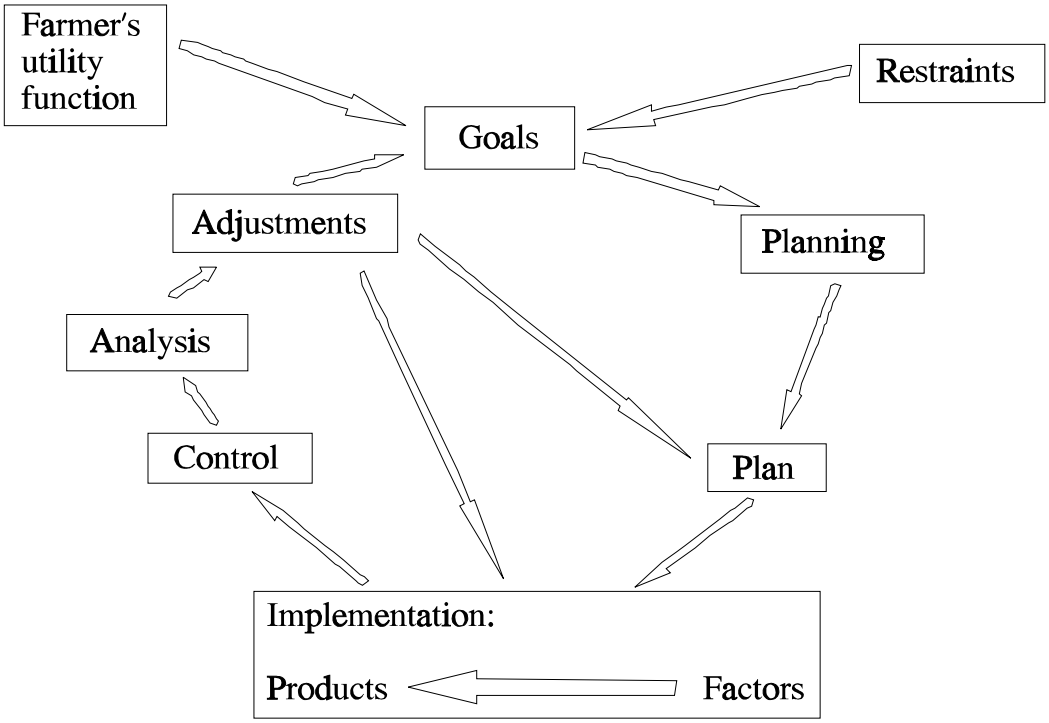

Figure 6. The elements of the management cycle.

instance he may give higher priority to leisure time or working conditions as he becomes older). Goals on the other hand may change as the conditions change. They are derived from the farmer's preferences in combination with the restraints, and since we noted in Chapter 1 that for instance legal restraints may very well change over time, the same of course apply to goals. The purpose of goals is only to set up some targets which (if they are met) ensure that under the circumstances defined by farmer's preferences and the restraints the production will be successful. In practice, goals may be expressed as a certain level of production, a certain efficiency etc.

It should be noticed, that goals are often defined as a result of planning under a longer time horizon than the one considered. Traditionally three different time horizons (levels) are distinguished. The *strategic* level refer to a long time horizon (several years), the *tactical* level refer to an intermediate time horizon (from a few months to a few years) and the *operational* level refer to a short time horizon (days or weeks). Thus goals for the operational level are typically defined at the tactical level.

When the goals have been defined the process of *planning* may be initiated. The result of the process is of course a *plan* for the production. A plan is a set of decided actions each concerning the future allotment of one or more factors. Alternative actions are evaluated on their expected utility as discussed in Chapter 2. Accordingly, the expected resulting production from each plan has to be known in order to be able to evaluate the utility (cf. Eq. (6)). So, what the plan actually contains is a detailed description of the factor allotment and the expected resulting production.

The next element of the management cycle is *implementation*. From a theoretical point of view this element is trivial (but certainly not from a practical). Implementation is just to carry out the actions described in the plan, and during the production process the factors are transformed into products.

During the production process, some *registrations* are performed. The registrations may refer to factors as well as products. Based on the registrations, some *key figures* (e.g. average number of piglets per farrowing, average milk yield per cow) describing the performance of production may be calculated. During the *control* process these calculated key figures are compared to the corresponding expected values known from the planning process.

The result of the comparison may either be that the production has passed off as expected or that one ore more deviations are identified. In the first case, the production process is continued according to the plan. In the latter case, the deviations have to be *analyzed*. The purpose of the analysis is to investigate whether the deviations are significant from (a) a statistical point of view **and** (b) from a utility (often economic) point of view. Because of the random elements of the production function ($\mathbf{e}_t$ of Eq. (1)) and because of observation errors relating to the method of measurement it may very well happen that even a considerable observed deviation from a statistical point of view is insignificant. That depends on the magnitudes of the random elements and the sample size (for instance the number of animals). Even if a deviation is significant from a statistical point (because of small random variation and a large sample) it may still turn out to be insignificant from a utility point of view. In the following section the statistical analysis is discussed in more details.

If it is concluded during the analysis that a deviation is significant from a statistical point of view as well as a utility point of view, some kind of adjustment is necessary. Depending on the nature of the deviation the adjustment may refer to the goals, the plan or the implementation. If the deviation concerns the factor allotment, the implementation has to be adjusted. During the planning process certain factor levels were assumed, but during the control it appeared that the *actual* factor allotment was different. Accordingly, something went wrong during the implementation of the plan.

If the deviation concerns the output level (i.e. the products), the conclusion depends on whether or not a deviation concerning the factor levels was found simultaneously. In that case, the deviation in output level is probably only a result of the deviation in input level. Accordingly the adjustments should focus on the implementation process.

If, however, there is a deviation concerning output, but none concerning input, we really face a problem. During the planning process, we assumed that if we used the factors represented by the vector **x** we could expect the production E(**Y**). Now, the control process show that the actual factor allotment *was* **x**, but the actual production was **Y'** which differs significantly from E(**Y**) both from a statistical and a utility point of view. If we consult Eq. (1), we have to conclude, that the only possible explanation is that we have used a deficient production function $f_{s,i}$ (in other words, the *validation* of the model used has been insufficient). The reason may be that the state *i* of the production unit *s* differed from what we assumed during the planning process, or it may simply be because of lacking knowledge on the true course of the production function. Under all circumstances, the situation calls for a new planning process where correct assumptions are made. Accordingly the adjustments refer to the plan.

During the new planning process it may also turn up, that one or more goals are impossible to meet, and in that case they have to be adjusted.

Finally, it should be emphasized that if any restraints (legal, economic, physical or personal) changes, new goals have to be defined and new plans have to be made.

3.2. Statistical evaluation of deviations identified during the control process

Whereas the analysis of a deviation from a utility point of view depends very much on the preferences of the individual farmer, the anaslysis from a statistical point of view is more general. In this section we shall briefly discuss some principles of this evaluation.

As mentioned in the previous section, the random variation associated with a calculated key figure has two sources: The observation error and the sample error. We shall denote the true (but unknown) value of the key figure as k_t and the calculated value as k_c. The relation between the two values may be expressed as

$$k_c = k_t + e_s + e_o \,, \qquad \textbf{(29)}$$

where e_s is the sample error and e_o is the observation error. Depending on the trait measured by the key figure one or both of the error terms may be zero or at least insignificant in magnitude. In the following, we shall investigate the consequences of (29) in relation to some practical examples.

3.2.1. Example 1: Milk yield of dairy cows

In the Danish management information system for dairy cows, the average daily milk yield of the cows during the first 24 weeks of lactation is calculated for first lactation cows and other cows, respectively. In the system, the farmer may compare the results to his own goals (expected values). The results in a herd were as shown in Table 3. We observe, that the calculated result for all parities is lower than expected. The question is now, whether or not the deviation is significant from a statistical point of view.

Table **3**. Average daily milk yield (ECM) during the first 24 weeks of the lactation in a herd.

Parity	Goal (expected value)	Calculated result	No. of cows	Acceptable result
1	23.5	20.4	10	21.7
≥ 2	27.8	26.0	12	25.6

According to Kristensen (1986) the standard deviation of **cumulated** milk yield (first 24 weeks) between cows in the same herd is 490 kg milk for first lactation and 630 kg milk for other lactations. Since these standard deviations have been estimated from exatly the same kind of basic registrations as those of the management information system they actually represent the standard deviation of the **sum** of the error terms (i.e. $e_s + e_o$). In order to use them in relation to the results of Table 3 they have to be corrected to daily levels. Since 24

weeks are equal to 168 days, the standard deviation of daily milk yield between cows is 490/168 = 2.92 kg for first lactation cows and 630/168 = 3.75 for others. Since the calculated results are average values of 10 and 12 observations, respectively, the standard deviations of $e_s + e_o$ become $2.92 \times 10^{-½} = 0.92$ and $3.75 \times 12^{-½} = 1.08$. As a rule of thumb for normally distributed data, a deviation is significant if it exceeds the standard deviation multiplied by 2, so in this case, acceptable results would be $23.5 - 2 \times 0.92 = 21.7$ and $27.8 - 2 \times 1.08 = 25.6$ as indicated in the last column of Table 3.

Since the calculated result concerning first lactation cows is clearly below the deduced acceptable value, we conclude that the deviation is significant from a statistical point of view. On the other hand, the result concerning other cows is higher than the acceptable value, and accordingly we have to accept the result as a possible consequence of usual random variation. In other words, the result concerning those cows does **not** call for any adjustments.

3.2.2. Example 2: Daily gain for bull calves (or slaugther pigs)

Usually, total gain y_t for bull calves or slaugther pigs over a certain period is calculated as follows:

$$y_t = x_1 + x_2 - x_3 - x_4 , \tag{30}$$

where x_1 is the total live weight of all calves delivered during the period, x_2 is the valuation weight (total weight of all calves present in the herd) at the end of the period, x_3 is the valuation weight at the beginning of the period and x_4 is the total weight of all calves born/inserted in the herd during the period. In order to arrive at the *daily* gain y_d, the total gain y_t has to be divided by the total number of days in feed z, i.e. $y_d = y_t / z$, where days in feed are calculated as

$$z = \sum_{j \in H} d_j , \tag{31}$$

where H is the set of **all** calves that have been in the herd during the period (or just a part of the period), and d_j is the number of days the jth animal has been in the herd. We shall assume that the standard deviation in daily gain between calves of the same herd amounts to 200 g. The standard deviation of the sample error e_s (i.e. the standard deviation of the average result in a herd with 100 animals) is therefore $200 \times 100^{-½} = 20$ g.

Based on the feeding and the housing conditions the goal for daily gain has been set to 1150 g in a herd with approximately 100 bull calves. The actual results registered in the herd for a 3 month period appear from Table 4.

The standard deviation of the observation error e_o depends on the method used for registration of the individual results of Table 4. We shall assume that x_1 has been calculated from actual weighings of all calves delivered to the slaughter house. Accordingly, we assume that no observation error is associated with this figure. The uncertainty associated with the weight of the calves inserted will also be insignificant. What is left is only the valuation weight of all animals at the beginning and the end of the period, respectively. The observation error on these figures depends heavily on the method used. We shall assume the standard deviation of

Table **4**. Results used for calculation of average daily gain in a herd.

Total live weight of 25 calves delivered, kg	x_1	12,500
Valuation weight at the end of the period, kg	x_2	21,159
Valuation weight at the beginning of the period, kg	$-x_3$	−23,441
Total weight of 20 calves inserted, kg	$-x_4$	−950
Total gain during the period, kg	y_t	9,268
Total days in feed	z	8,827
Daily gain, g	y_d	1,050

the observation error (on x_2 and x_3) it is zero if all animals (approximately 100) have actually been weighed. If the total weight is calculated by weighing only representative calves, the corresponding standard deviation is assumed to be 5%. Finally, if the valuation weights are just assessed by visual observations, the standard deviation is assumed to be 10%.

Based on these assumptions the variance and standard deviation of the observation error may be calculated by standard methods as shown in Table 5, where also the total standard deviations of the herd result under the three valuation methods appear.

Table **5**. Error terms to be used in the calculation of the standard deviation of herd result.

Error term	Valuation method		
	Weighing[1]	Sample[2]	Visual
Variance of observation error associated with y_t	0	2,292,959	9,971,838
Standard deviation of observation error, y_t	0	1579	3158
Standard deviation of e_o (observation error, y_d)	0	179	358
Standard deviation of e_s (sample error)	20	20	20
Variance of $e_s + e_o$	400	32,441	128,564
Standard deviation of $e_s + e_o$, g	20	180	359

[1] Weighing of *all* calves at valuation.
[2] Weighing of a representative sample.

If, again, we use the rule of thumb, that a deviation has to be larger than twice the standard deviation in order to be significant, we conclude that acceptable results under the three valuation methods are 1150 − 2 × 20 = 1110 g, 1150 − 2 × 180 = 790 g and 1150 − 2 × 359 = 432 g, respectively. Accordingly, the calculated result of 1050 g is only significant from a statistical point of view if the valuation method has involved weighing of *all* calves.

3.2.3. Example 3: Reproduction in a dairy herd

The previous examples referred to variables which might be regarded as at least approximately normally distributed. We now turn to examples involving categorical data, where other distributions are involved.

If N dairy cows are inseminated, the number of resulting pregnancies n will be binomially distributed with parameters N and p, where p is the conception rate. In a dairy herd, the goal concerning the conception rate has been set to 0.5, wheras the actually calculated conception rate amounts to only 0.4. If we assume, that the state of a cow (pregnant or not pregnant) may be determined with certainty (i.e. no observation error), the result may be evaluated by a binomial test using the hypothesis p = 0.5. The test variable t_N is defined as:

$$t_N = P(n<n_o|p=0.5) \ , \qquad \textbf{(32)}$$

where n_o is the number of pregnancies observed. Whether or not the observed conception rate differs significantly from the goal depends heavily on the number of inseminations behind the calculation. In Table 6, the test variable t_N has been calculated under assumption of three different sample sizes.

Table **6**. Test of the hypothesis that the true conception rate is 0.5 given an observed rate of 0.4 and the sample size (number of inseminations) N.

Number of inseminations, N	Number of observed pregnancies, n_o	Test variable, t_N
10	4	0.17
20	8	0.13
50	20	0.06

As it appears from the table, the deviation is only close to be significant if the sample size is 50. If this figure is compared to usual herd sizes in European countries, we may conclude that the conception rate has to be calculated over a rather long period if shall pay any attention to deviations from the goal.

3.2.4. Example 4: Diseases

If the herd size and the disease incidence is constant, the number of cases of a specific disease may be represented by a Poisson distribution provided that individual cases appear independently of each others (which is *not* the case with contagious diseases). The actual number of cases may therefore be compared to the goal (or expected result) using probabilities calculated from a Poisson distribution taking the defined goal as it's parameter. The test variable t_N is calculated as

$$t_N = P(n \geq n_o | \lambda = goal) \ , \qquad \textbf{(33)}$$

where λ is the parameter of the Poisson distribution and n_o is the number of disease cases observed.

3.2.5. Concluding remarks

It should be emphasizes that the kind of statistical "tests" illustrated in the previous sections are only of indicative nature. They may not be confused with tests performed on data from controlled experiments. The purpose is only to provide a rough estimate of the significance of an observed deviation. The message is, that a result calculated from actual registrations is often associated with a rather big uncertainty as illustrated in the examples. In many cases it is possible to "estimate" that uncertainty rather precisely, but there is no correct level of significance to use in the evaluation of results. In a true situation the relevant level will also depend on the significance from a utility point of view.

An other important observation is that in order to be able to evaluate the production we have to perform some registrations relating to the production process. In an other note of this series the process of transforming such registrations to useable information is discussed.

3.3. Evaluation of deviations from a utility point of view

The evaluation of deviations from a utility point of view is just as important as the statistical analysis. If may very well happen that a deviation is significant from a statistical point of view (where significance is merely a question of sample size), but insignificant from a utility point of view (and vice versa). Adjustments are of course only relevant if the deviations observed are significant from both points of view.

The evaluation from a utility point of view is complicated by the dynamics of production. If, for instance, the average number of days open in a dairy herd is higher than expected (and the statistical analysis showed that it was a significant deviation) we have to consider what the deviation means for the farmer’s utility. The direct consequences for production of a cow conceiving 16 weeks after calving (b) instead of 12 weeks after calving (a) are illustrated in Figure 7.

As it appears from the figure, the direct consequences include:

a. Next calving is delayed by 4 weeks.
b. The milk yield towards the end of the lactation is slightly higher.
c. The number of days in milk is increased by 4 weeks.
d. The milk lactation curve of next lactation is shifted 4 weeks to the right.

Expressed numerically, the consequences might for instance be as indicated in Table 7, where the economical net returns are calculated initially on a lactation basis and afterwards normalized to annual figures per cow.

The economic value per cow per year of a decrease in the number of days open from 16 to

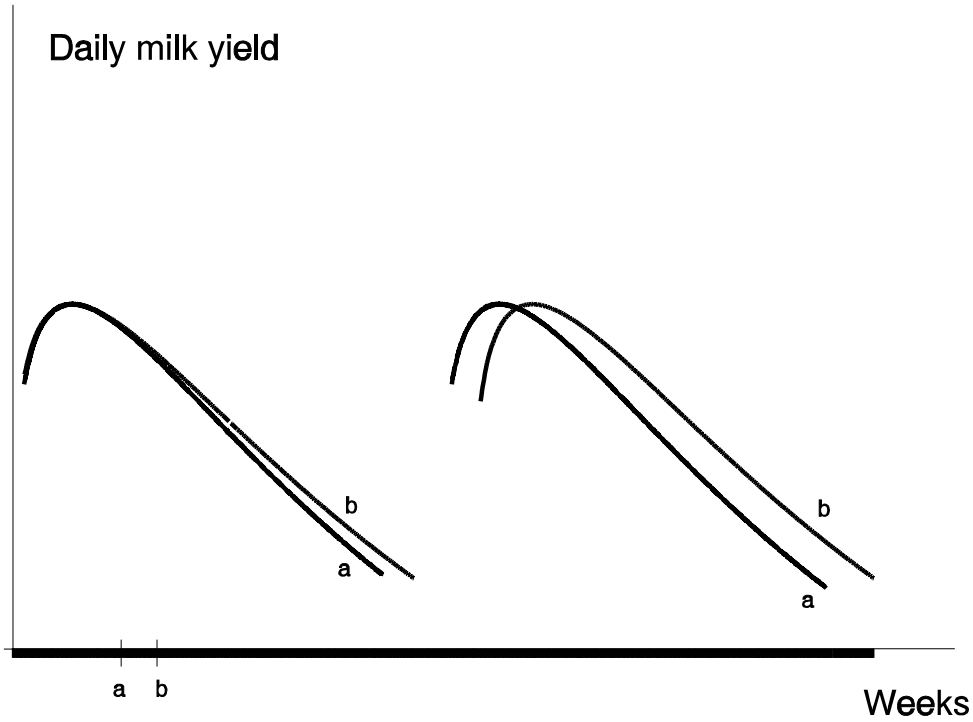

Figure 7. Daily milk yield during two lactations for a dairy cow conceiving 12 weeks after calving (a) compared to a cow conceiving 16 weeks after calving (b).

Table **7**. Comparison of lactation results for a cow conceiving after 12 weeks with those of a cow conceiving after 16 weeks.

	Price	Conception: 12 weeks		Conception: 16 weeks	
		Amount	Value	Amount	Value
Total milk yield, lacation	2.10[a]	6000	12600	6200	13020
Number of calves born	1000	1	1000	1	1000
Total economic value, lactation	-	-	13600	-	14020
Calving interval, days	-	364	-	392	-
Total economic value, year	-	-	13637	-	13054

[a] Marginal value: Price of milk less marginal feed costs.

12 weeks appears to be 13,637−13,054 = 584. The *annual* difference in number of days open

per cow is 16×7×365/364 − 12×7×365/392 = 20 days. Assuming linearity, the cost of one additional day open is therefore 584/20 = 29.

Assume now, that in a herd with 100 cows, the average number of days open has been calculated to 105 wheras the goal is 90. It would be very tempating to claim that the total economic loss caused by this deviation is (105-90)×29×100 = 43,500, but nevertheless, it would be a serious mistake. In order to reveal why, we shall consider the implicite assumptions behind the calculations above. Those assumptions are:

- **The *marginal* value of an improvement is assumed to be constant (i.e. independent of the initial level)**. In other words, a decrease from 200 days open to 199 days is assumed to be of the same value as a decrease from 30 to 29 days. Such an assumption is certainly *not* realistic.
- **All cows are assumed to be identical in the sense that they all conceive exactly after 12 (repectively 16) weeks.** The real situation is very different. Even though an average result is for instance 105 days, *individual* results vary considerably among the cows of a herd. Thus, if the assumption of constant marginal value does not hold, we need to know the distribution of individual results in order to be able to estimate the economic consequenses at herd level.
- **All cows are assumed to be identical concerning milk yield**. From the calculations it is obvious that a different level of milk yield would give us an other result.
- **Conception is assumed to be independent of milk yield**. In the real world, it is a well known problem that the conception rate of a high yielding cow is lower than that of a low yielding. It is therefore reasonable to assume that in particular the high yielding cows are those that cause a high average number of days open.
- **It is assumed that no cows are replaced**. In practice annual replacement rates of 30 to 50% are usual. It is obvious, that if a cow is replaced, the effects of delayed calving and delayed new lactation are irrellevant.
- **The farmer's management is assumed not to interact with the observed result.** The farmer probably has some policy concerning for instance replacement. If a cow is high yielding, he accepts more days open than if it is low yielding. If the number of days open grows, he probably also increases hes replacement rate.
- **It is assumed that there is no milk quota.** In case of a milk quota, the marginal value of an increased milk yield is much lower than assumed in the example. The reason is that if a higher milk yield is obtained (for instance because of improvements in reproduction) the farmer has to decrease the number of cows in order not to violate the quota (refer to Kristensen & Thysen, 1991, for a discussion). Presence of a milk quota may reduce the costs of an additional day open by 60-70% as shown by Sørensen & Enevoldsen (1991)

These considerations clearly illustrate that calculations as those of the example are of little value in the evaluation of results obtained in production. It is very important that random variation, interactions, management policies and restraints are taken into account in such calculations. The most obvious tool to apply is certainly simulation, but also some of the other tools discussed in Chapter 4 may be used in some cases. Evaluations from a utility point of view remains, neverheless a true art.

4. Decisions and strategies: Framework and techniques

In this chapter we shall initially discuss general aspects of decision making in animal production with emphasis on information needs, time horizons and production units. Following that discussion, a general overview of techniques available in the determination of optimal decisions and strategies is given.

4.1. The framework of decision making in animal production

During the planning process we have to make decisions. We shall define a decision as the allotment of some factor used in the production. The purpose of planning is of course to make one or several *optimal* decisions so that the set of decisions making up the plan maximizes the expected utility relating to the relevant time horizon as discussed in Chapter 2.

4.1.1. Information needs

From Eq. (7), we observe that in order to be able to make an optimal decision, *in general*, the following information is necessary:

1. The state of the production unit at *all* stages of the planning horizon.
2. The production function(s) and the distribution(s) of the random term(s).
3. All attribute functions relevant to the farmer's preferences.
4. The utility function representing farmer's preferences.

Recalling the discussion of Chapter 1, we have to add the following information:

5. All restraints of legal, economic, physical or personal kind.

As concerns item 1, it is obvious, that the future state of the production unit is not always known at the time of planning. The state represents all relevant information on the unit (i.e. the traits of the unit) and very often it varies at random over time. It is therefore likely, that it is not possible to make decisions relating to all stages at the same time. Instead, we may choose a *strategy* (or *policy* as it is also some times denoted) which is a map (or function) assigning to each possible state a decision. In other words, having chosen a strategy we may at all stages observe the state of the unit and make the decision provided by the strategy.

The identification of an optimal strategy is complicated by the fact that decisions may also influence the future state of the unit as illustrated by the following example. Let us assume that the unit is an animal, that the relevant information (the state) is the weight i of the animal and what we have to decide is the feeding level x of the animal. It is obvios, that the optimal level of x will depend on the weight of the animal i, but it is also obvious, that the feeding level x will influence the future weight j of the animal.

A thorough inspection of Eq. (7) will unveil that the decision made may also influence the *future* production, but, since the state i is assumed to contain all relevant information on the production unit, we may just include information on previous decisions in the definition of the state. Thus we do not need to consider that aspect further. Consider again the example of

Chapter 1, where we noted that the feeding level of a heifer influences the future milk production as a dairy cow. If we denote the feeding level at various stages as x_1, x_2,... , the stage of first calving as k and the milk yield of stage n as y_n we might express the production function in the same manner as in Eq. (1):

$$y_{k+n} = f_{si}(x_{k+n},...,x_1) + e_{k+n} , \quad \textbf{(34)}$$

where the state i for instance may be defined from combinations of traits like lactation number, stage of lactation and body weight. Thus the state space Ω_1 is defined as the set of all possible combinations of those individual traits, and an individual state $i \in \Omega_1$ is any combination of values representing the three traits. If, however, we redefine the state space to include also the feeding levels of previous stages (i.e. $\Omega_2 = \Omega_1 \cup \{x_1,...,x_{k+n-1}\}$) exactly the same information may be contained in the relation

$$y_{k+n} = f_{si}(x_{k+n}) + e_{k+1} , \quad \textbf{(35)}$$

where, now, $i \in \Omega_2$.

Eq. (35) illustrates that the state concept is very essential for the planning process. The identification of the state of the production unit is therefore a very important task. In order to be able to identify the state we have to perform some registrations concerning the production unit. It may some times happen that we are not able to observe the traits of the state space directly, but only indirectly through other related traits. The precision of our knowledge concerning the true state may therefore vary, but Bayesian updating methods are available for handeling such situations with imperfect knowledge. Those aspects are dealth with in a later chapter.

We may now reformulate the information needs when choosing an optimal strategy:

1. A production function describing the immediate production given stage, state and decision and the distribution of the possible random term(s).
2. The distribution of the future state given stage, state and decision.
3. All attribute functions relevant to the farmer's preferences.
4. The utility function representing farmer's preferences.
5. All restraints of legal, economic, physical or personal kind.

As concerns items 3 and 4, the discussion of Section 2.2.3. should be remembered. Several techniques are available for identification of optimal strategies.

4.1.2. Levels

In herd management we face a hiearchy of decisions made at different levels with different time horizons. In this section, and the following section, we shall consider the implications of this hierarchy.

By level we mean the production unit considered. As discussed in Chapter 1 the unit may be an individual animal, a group or a pen, a section or even the entire herd. For instance the decision to build a new barn is an example at herd level. Decisions concerning feeding are

typically made at group or section level wheras decisions concerning culling or medical treatment may be made at animal level.

As long as decisions at different levels are mutually independent we may solve a problem at one level without considering the other level. Unfortunately it is very rare that such independence exists as the following example shall illustrate.

We want to build a new barn for our dairy cows, and two alternatives, a_1 and a_2, are available. The two barns of course differ in several respects, but one of the differences is that barn a_1 allows for grouping of the cows wheras grouping is not possible in barn a_2. This difference means that in order to make an optimal decision concerning the kind of barn to build, we also have to consider how to feed the animals in each of the two systems, because the option of grouping makes other feeding strategies possible than if all cows are housed in a single group. It then turns out that feeding strategy α_1 is optimal in barn a_1 whereas strategy α_2 is optimal in barn a_2.

Unless α_1 has exactly the same utility consequences as α_2, the differences must be taken into acount when the optimal barn is chosen.

4.1.3. Time horizons

A general aspect of herd management under risk is that decisions have to be made without certainty about the future state of the production unit. The uncertainty increases with the time horizon of the decision, i.e. it is more prevalent at the tactical level than at the operational level. Having made a decision at the tactical level, the manager is restricted by the consequences for the duration of the time horizon. It may very well later turn out, that the actual state of the production unit differs from the expected state at the time of the decision, but the only way the manager may adjust to the new situation is by making decisions at the operational level. These decisions should be conditionally optimal given the tactical decision made and the current state of the production unit. In other words, the decisions at the operational level may be regarded as a way of adjustment to risk and in that way compensate for the incomplete knowledge on the future state of the production unit.

In general, it must be assumed that if decision a_1 is made at the tactical level, then strategy α_1 is optimal for decisions at the operational level (a strategy is defined as a set of decisions relating to the set of possible states of the production unit). On the other hand, if decision a_2 is made at the tactical level, then strategy α_2 is optimal at the operational level. It will be an exception, if $\alpha_1 = \alpha_2$. In other words, it is not possible to choose an optimal decision a' at the tactical level, unless a conditionally optimal strategy α' has been determined at the operational level (conditional given the tactical decision).

In case of a management problem with limited time horizon (for instance the duration of the tactical decision considered) the mutual dependency between decisions at the tactical and operational level is not really a problem. We just have to determine optimal policies at the operational level given each of the alternative tactical decisions and, afterwards, to choose the tactical decision maximizing the objective function. A problem corresponding to this situation is discussed by Jensen (1995), who considered optimal mildew management policies in winter wheat under different nitrogen fertilization strategies. In that example, the decision at the

tactical level is to choose a nitrogen fertilization strategy and the decisions at the operational level are to treat the crop for mildew. The time horizon is limited to the growing season just like the tactical decision. The problem was solved within the framework of a Bayesian network in combination with a usual backwards dynamic programming algorithm.

If, however, the time horizon is unknown or at least very long as it is typically the case in animal production, the situation is far more complicated. Examples of tactical decisions include mating of a female animal with a male animal of a certain quality or choosing a certain feeding level for an animal. Such decisions have (depending on the animal species and other circumstances) a time horizon of a few months, but unlike the mildew management problem, the time horizon of the production is not limited to a growing season or the like. Instead the production is continuous, which is often modelled by an infinite time horizon. In order to cope with such a situation, the decisions at the tactical and operational level have to be optimized simultaneously in order to ensure over-all optimality.

The terms "tactical" and "operational" are of course rather arbitrary. In general we have to deal with decisions at several levels having different time horizons.

4.2. Methods

The techniques will be described according to their ability to represent the various kinds of necessary information discussed in Chapter 4. Furthermore, their potentials for integration of decisions at different levels and time horizons are discussed. It is not the purpose to describe the various methods in details, but only to provide a general survey relating to the issues of the previous chapters. For each method, references to relevant literature is given.

4.2.1. Rule based expert systems

Research concerning expert system is a recent development within the area called Artificial Intelligence (AI). The British Computer society has defined expert systems as follows: *An expert system is regarded as the embodiment within a computer of a knowledge-based component from an expert skill in such a form that the system can offer intelligent advice or take intelligent decision about a processing function. A desirable additional characteristic, which many would consider fundamental, is the capability of the system, on demand to justify its own line of reasoning in a manner directly intelligible to the enquirer. The style adopted to attain these characteristics is rule-based programming. (Cf Dindorph, 1992)*. This is just one of many proposed definitions of expert systems.

The fundamental difference between the rule based systems, and the approach that we have presented until now, is that instead of trying to model a system, the rule based expert systems tries to model the expert, or rather the expert’s approach to problem solving. Originally, the ambition within AI-research was to make general problem solvers that could be used for any problem area, but this was realized by most researchers within the area to be overambitious. The research efforts had, however, led to new approaches towards problem solving, and within narrow problem (expert) domains, the approach showed some promising results.

Rule based expert systems have three components: a knowledge base containing the expert’s knowledge of the domain, an inference engine that decides how and when to use the

knowledge, and a user interface.

The knowledge base contains knowledge of a problem domain, as it is described in text books, as well as and expert knowledge, e.g. exceptions to general rules, experiences from previous problems, and time-efficient approaches on how to solve problems within the area. The knowledge base is an enhanced data base that apart from data also contains logic rules for the connection between the items in the knowledge base, e.g *if x or y then z*.

The inference machine contains the mechanism for deduction based on the logical rules in the knowledge base. The deduction can use different inference principles, such as backward chaining and forward chaining. In the rule showed above, the inference machine would start out by finding the value of *z*, and given the knowledge of *z* establish the value of *x* and *y*. Forward chaining would start out by establishing the values of *x* and *y* and deduce the value of *z* subsequently. In both cases the unknown values are found either by a question to the user or by combining other rules in the knowledge-base. The optimal choice of inference principle depends on the type of problem the expert system is supposed to solve. In very complex expert systems, neither forward- or backward chaining is fast enough, and the so-called heuristic search strategies are needed. These strategies work primarily by searching the knowledge base in an efficient order, focusing on areas, where a solution to the problem is most likely to be found. Both general and problem specific heuristic strategies exist.

The questioning mechanism is a standard part of the user interface. Besides posing questions, the user interface is usually able to explain, why it asks the question, i.e. 'I'm trying to establish the value of *z* because I want to know if either *x* or *y* is true. Another facility is the explanation facility, i.e. 'I know that neither *x* is true nor *y* is true because *z* is not true. Usually the phrases are formulated more user friendly.

Programs for maintaining knowledge bases in connection with inference machine and user interface as an integral part is sold as the so-called expert system shells. This concept originates from the medical diagnosis system, MYCIN. The knowledge base in MYCIN was emptied and the program sold as E-MYCIN (or empty mycin), and was thought to be applicable to any problem domain. Very often these shells are programmed in programming languages where logic deduction can easily be represented, such as LISP or PROLOG, but standard programming languages such as C and Pascal can of course be used.

Rule based expert system can be categorized into several areas (Hayes-Roth *et al.*, 1996). Referring to Figure 6 they comprise the planning, control and analysis phase of the production cycle.

In developing rule based expert systems two 'players' are essential, of course the expert, but in addition the so-called knowledge engineer. The role of the knowledge engineer is to 'extract' the knowledge from the expert and to formulate the knowledge as rules that can serve as input to the knowledge base. Knowledge engineering has in fact become a research area in its own right.

To illustrate the problem of knowledge engineering, the first rule based expert systems were based on very simple 'expert' rules very much inspired by the diagnostic systems, e.g. if indication a and b are observed then problem is probably c. Later on it was realized that the

expert relies on many information sources and part of being an expert is to know when to draw on which knowledge sources. If we look at the information necessary to make optimal decisions as mentioned in section 4.1.1 this can be seen as the result of an expert's problem solving. An expert system would therefore guide the user through obtaining the necessary information. If it is not possible to obtain the necessary information it would use other information and based on the expert's experience try to make a sufficiently good plan.

The current trend is that the rule-based system does not function as standalone systems, but rather as an integral part of other systems, the so-called knowledge based systems. A typical example could be that the expert system helps in establishing the user's utility function by asking questions and then uses this utility function when calling an optimizing program. The concept is now incorporated into the wizards and experts known from standard computer program, e.g. spreadsheets and word processors.

4.2.2. Linear programming with extensions

The general linear programming problem may in matrix notation be written as follows:

$$\begin{aligned} &\mathbf{px} = \mathbf{Min!} \\ &\textit{subject to} \\ &\mathbf{Ax} \leq \mathbf{b} \\ &\mathbf{x} \geq 0 , \end{aligned} \tag{36}$$

where **p** is a constant row vector with m elements, **A** is a constant matrix in $n \times m$ dimensions, **b** is a column vector with n elements, and **x** is a vector of variables. The problem is to select a vector **x** that minimizes the linear objective function **px** and simultaneously meets the linear restraints $\mathbf{Ax} \leq \mathbf{b}$ and $\mathbf{x} \geq \mathbf{0}$.

Eq. (36) represents the standard formulation of a linear programming problem. In applied models it is often convenient to define a maximization problem instead, and some of the restraints may be of the kind $\mathbf{a}_i\mathbf{x} \geq b_i$ or the kind $\mathbf{a}_i\mathbf{x} = b_i$, but any linear programming problem may be rearranged in accordance with the standard formulation of Eq. (36).

If we interpret the linear programming problem in relation to a herd management decision problem, then **x** is a vector of factor levels and $\mathbf{Ax} \leq \mathbf{b}$ is a set of restraints of legal, economic, physical or personal kind. It should particularly be noticed, that personal restraints may also include restraints on levels of attribute functions (for instance leisure time or monetary gain). This direct representation of restraints is probably the main force of the method. The objective function **px** has to represent the aggregate utility function.

If we compare the linear programming problem with the information needs of a decision problem (cf. Chapter 4) we observe that all random elements are missing. At least in the standard formulation, the method is deterministic. Also the dynamic linking to the future state of the production unit is missing. A consequence of the latter shortcoming is that only effects at the current stage are represented. In other words, the method is static of nature. Furthermore, we observe, that since the aggregate utility function has to be linear in the factor levels **x**, also the production function, all attribute functions and the utility function have to

be linear. Examples of linear attibute functions are shown as Eqs. (8), (11) and (14), and a linear utility function is shown as Eq. (24). This demand for linear functions and linear restraints is a serious weakness of the method.

Several of the shortcomings mentioned may be redressed or at least adapted by extensions to the method: The linear objective function may be replaced by a quadratic one (quadratic programming); the static nature may be modified by introduction of stages and additional restraints ensuring dynamic links (dynamic linear programming); random terms may be added to the elements of **A** and **b**, and the corresponding restraints may be expressed as probabilities (stochastic programming); and often, non-linear functions may be approximated by pieces of linear relations over short intervals. In particular, dynamic linear programming, may be used to link decisions at different levels with different time horizons.

Herd management applications of linear programming are numerous. The most frequent application is no doubt for ration formulation, where least-cost rations meeting the nutrinional requirements of the animals in question are met. Most often such programs ignore the effect of feeding on production. An example of such a ration formulation program developed for teaching has been described by Kristensen (1993)[1].

Also examples of application of linear programming for whole-farm planning may be found in literature. Refer for instance to Hansen (1992) and Hardie (1996). Such models are often very large containing thousands of variables and restraints, but since very efficient standard software is available this is hardly a problem.

4.2.3. Dynamic programming and Markov decision processes

Consider a production unit which is observed over a number of stages $n = 1,...,N$. At the beginning of each stage, we observe the state $i \in \omega_n$ of the unit. Having observed the state, we have to take an action $d \in D_n$ concerning the production unit. Usually, the state space ω_n and the action space D_n are assumed to be finite sets. Depending on the stage, state and action, a reward is gained. The reward may very well be a random variable, but the expected value $r_i^d(n)$ has to be known. Also the state to be observed at the next stage is a random variable. We shall denote as $p_{ij}^d(n)$ the conditional probability of observing state i at stage n+1 given that state i has been observed and action d taken at stage n. Finally, a strategy s is defined as a map assigning to each combination of stage and state an action $s(n,i) \in D_n$.

The purpose of dynamic programming is to determine a strategy which (in some sense) is optimal. Several optimization techniques are available. The most commonly applied method is called *value iteration* where a value function representing the expected total rewards from the present stage until the end of the planning horizon (i.e. stage N) is maximized. Optimal decisions depending on stage and state are determined backwards step by step as those maximizing the value function. This way of determining an optimal *policy* is based on the Bellman principle of optimality which says: "An optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal

[1] The program itself is available as a Zip-file on www at the following URL: http://www.prodstyr.husdyr.kvl.dk/Software/ah.html

policy with regard to the state resulting from the first decision" (Bellman, 1957 p. 83). Value iteration is often just denoted as *dynamic programming*.

If N is large, an infinite planning horizon is often assumed. A relevant optimization technique for infinite stage problems is *policy iteration*. This method was introduced by Howard (1960), who combined the dynamic programming technique with the mathematically well established notion of a *Markov chain*. A natural consequence of the combination was to use the term *Markov decision process* to describe the notion. The policy iteration method was a result of the application of the Markov chain environment and it was an important contribution to the development of optimization techniques.

The objective function being maximized during optimization depends on the circumstances. It may represent the total expected rewards, the total expected *discounted* rewards, the average rewards per stage or the average rewards over some kind of physical output.

If we compare the dynamic programming problem with the information needs of a decision problem (cf. Chapter 4) we observe that most aspects are covered. The rewards directly correspond to production functions; the conditional probabilities $p_{ij}^{d}(n)$ represent the dynamic random links to future stages; and the objective function represents the farmer's utility function. There are, however, some restrictions on the kind of utility function which can be represented in a dynamic programming model. The restrictions concern the way in which individual stage attributes are aggregated into the over-all utility function as described by Eq. (5). In order to be able to use dynamic programming, we implicitely assume that the aggregation may be performed in such a way that we first aggregate attributes at the same stage into a stage specific utility v_n of the kind

$$v_n = g_n(u_{1,t_n},\ldots,u_{k,t_n}) \ , \qquad \textbf{(37)}$$

where g_n is a stage specific utility function of arbitrary kind. In the dynamic programming context, v_n is identical to the reward $r_i^d(n)$. The over-all utility (i.e. aggregation over stages) in turn must be calculated as a simple sum of the stage specific utilities v_n, as a discounted sum (cf. Eq. (9)), as the average value over stages *or* as the average value over some kind of physical output or input from production.

The most difficult kind of information to represent in dynamic programming models is the information on restraints. There is no general solution to that problem, but some times it may be solved by using an objective function maximizing average rewards relative to the most limiting restriction. An example is maximization of average net returns per unit of milk produced under a milk quota (Kristensen, 1989). In other cases combination of the method with methods like simulation (Ben-Ari & Gal, 1986; Kristensen, 1992) or genetic algorithms (Houben, 1995) may circumvent the restraint problem.

A major problem in relation to dynamic programming models is the so-called *curse of dimensionality*. Since the state space is represented by discrete levels of a set of traits (state variables), models tend to become very big. Thus a model presented by Houben et al. (1994) contained 6.8 million states. Despite the size of the model, optimization was still possible due to a new notion of a hierarchic Markov process described by Kristensen (1987, 1991).

Numerous applications of dynamic programming are described in literature. A relevant textbook concerning application in agriculture has been written by Kennedy (1986). The book also contains a survey of agricultural applications. In herd management, the technique has most often been applied for operational decisions like replacement, insemination and medical treatment of animals.

4.2.4. Probabilistic Expert systems

Another part of the research area named Artificial Intelligence are the so called probabilistic expert systems that rely on the Bayesian network. The following description is based on Lauritzen (1995).

In some areas where expert systems are appropriate, the task involves performing a sequence of steps according to specified logical rules. However, other expert systems work in domains that are characterized by inherent uncertainty. This uncertainty is either due to imperfect understanding of the domain, incomplete knowledge of the state of the domain at the time where the task has to be performed, randomness in the mechanisms governing the behaviour of the domain or a combination of these. Within these domains probability and statistics can serve to represent and manipulate the uncertain aspect of domains having these characteristics. Probabilistic methods were for some time discarded in this context as requiring too complex specification and computation. However, the work of Pearl (1988) and Lauritzen and Spiegelhalter (1988) demonstrated that these difficulties could be overcome, based on causal networks or as it is now usually termed Bayesian networks. There exist other formalisms for handling uncertainty in expert system, such as the fuzzy sets, but these will not be discussed in the present context.

The rule based systems were mainly constructed through modelling of the behaviour of the expert and the encoding this behaviour in terms of rules of various kind. In contrast, probabilistic expert systems are constructed by modelling the domain rather than the expert. The method is thus in close correspondence with the approach used in this book, where the domain is modelled using production functions etc. The probabilistic expert systems specify a graphical model for the variables. The reasoning is then performed by updating the probabilities of the domain in the light of the specific knowledge according to the laws of conditional probability.

The graphical model captures association between entities in the domain, or rather lack thereof, in terms of conditional independence that in a systematic fashion are encoded in a *graph* or *network* with *nodes* representing the entities themselves and *edges* representing associations between them. The *nodes* are represented as dots or circles. The *edges* are either *directed* corresponding to influences of a causal nature and represented as arrows, or *undirected* corresponding to symmetric associations (e.g. correlations) and represented as lines.

The use of the graphic specification in the probabilistic expert systems plays several roles. For example, it gives a visual picture of the domain information; it gives a concise presentation of domain information in terms of conditional independence relations, and it enables rapid computation and revision of interesting probabilities.

The graphic method can also be used for several important tasks in the specification process.

It can be used to *learn* quantitative and structural aspects, or as it known within general statistics, estimation and model selection.

If we compare the probabilistic expert systems method with the information needs of the decision problem, it is important to recognize that the method is inherently a static method, even though attempts have been made to model dynamic systems as well. For control and analysis purposes it is ideal, i.e. the method can assign probabilities to observed deviations, whether they are random or not. It can also make a diagnosis in the analysis, that is, indicate probabilities for different causes of the deviation. This can in turn serve as the necessary basis for decisions concerning changes in production plan.

The expert systems can be build as recurrent time slices and can in this manner represent dynamic production functions. And thus can predict the future state for given decisions. The restraints concerning the production function can be modelled, but restraints may cause the same problems as described under Dynamic Programming and Markov decision processes.

To represent decision in Bayesian networks the decision can be included as a random variable in the model, with the different decisions as level of the variables. When the decision is made, the corresponding level is assigned a probability of 1. This approach does not make any search for optimal decisions.

If decisions have to be included, *Influence diagrams* should be used instead. *Influence diagrams* can in fact be fitted into the general frame work of Bayesian Networks.

4.2.5. Influence diagrams

Influence diagrams was introduced by Howard and Matheson (1981) as a formalism to model decision problems with uncertainty for a single decision maker. The influence diagrams gave a more compact graphical representation of a decision problem than the more traditional decision tree approach. The influence diagram is very similar to the Bayesian network consisting of chance node and arrows denoting causal effect. In addition two more node types are introduced, the decision node shown as a square, and the value node shown as a diamond.

Originally, the influence diagrams was translated to a decision tree within the computer and the standard "average-out and fold-back" algorithm were applied on that tree. In Shacter (1986) a method was suggested for solving the decision problem represented by the influence diagram directly, without the translation to a decision tree. This method transformed the influence diagram by successively removing nodes in the graph, until at last only one final utility node remained, holding the utility of the optimal policy. In order to solve many similar problems one therefore had to start from scratch every time. The transactions performed on the diagram consisted of four simple transactions, the arc reversal (application of Bayes Theorem), node removal by summing out, expectation of a value node with respect to a change node, and finally removal of a decision node into a value node by maximization. Initially problems were formulated with only one value node. By introducing the concept of separability of the utility function or value function, Tatman and Shachter (1990) showed that the Dynamic Programming Problems could be solved within the Influence Diagram framework, by introducing the separability of the utility function. The requirement for separability is the same as the requirement that the overall utility is calculated as a simple

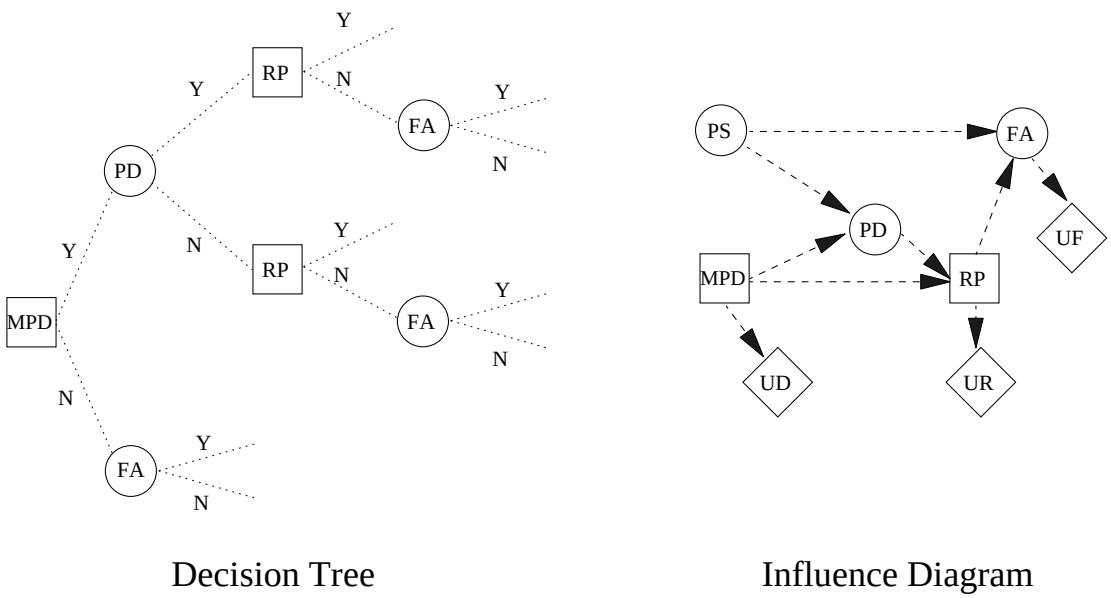

Figure 8: Decision tree and Influence diagram representation of the pregnancy diagnosis and replacement problem.

sum of stage specific responses as mentioned in section 4.2.3.

Shenoy (1992) proposed another algorithm that gave the solution to the influence diagram without disrupting the structure of the diagram. Then Jensen *et al.* (1994) showed how a similar approach could be incorporated within the general framework of Bayesian Networks. This approach has been implemented in the Hugin expert system shell. The similarity between influence diagrams and Bayesian networks means, that several important improvements is to be expected. This comprises, e.g approximate solutions by techniques such as Gibbs sampling (Bielza *et al.* 1996; Charnes and Shenoy, 1996), easy representation of dynamic models as in dHugin (Kjærulff, 1995) and easy calculation of promising strategies, e.g. the best N strategies (Nilsson, 1996).

As influence diagrams closely corresponds to Dynamic Programming the same comments concerning the information needs of the decision problems can be made. In addition, the current version of influence diagrams are inherently static, and no algorithm corresponding to policy iteration has been found (even though R.A. Howard's research has been central for both developments). If stages of varying time length have to be modelled, time has to be included in the model and the discounting factor has to be incorporated directly in the utility. Furthermore decision between qualitatively different subprocesses, such as in Hierarchic Markov processes is currently not possible. It should, however, be noted that it is a very active research area, and continuous progress is to be expected.

Thus, the method has obvious possibilities for application within animal production, but so

far the only example known to the author is an investigation of the value of using feed analysis for feed composition in dairy cattle (Pedersen, 1996). Within crop protection a system has been made for decision making concerning mildew management in winter wheat (Jensen, 1995).

4.2.6. Simulation

As the name implies, a simulation model is simply a model of a system. The model is used for the study of the real system's behaviour under different conditions. Within animal production the term usually refers to computer based dynamic calculation models.

Formally, the simulation model is a computer representation of the production function, the attribute function, and/or the utility function. The degree of detail differs between the different models.

The input to the model consists of two elements, a set of parameters, Φ, and a set of decision rules, Θ. The decision rules specify the setting of input factors as well as other decisions in the system. The term *decision rule* is used rather than *decision strategy*, because usually no direct mapping between the rule and the state of the whole system exists. A *decision rule* can e.g be to use a dynamic programming model to specify a decision strategy every (simulated) year. Another example is to use some simple rule-of-thumb (heuristic) to make culling decisions. The set of parameters can be split in two, $\Phi=(\Phi_0,\Phi_{s\cdot})$, where Φ_0 are the initial values of the parameters at the start of the calculation (State of Nature) and $\Phi_{s\cdot}$. is parameter values that change during simulations. The elements of $\Phi_{s\cdot}$ are often called state variables, and can be further split into time stages of the model i.e. $\Phi_{s\cdot}=(\Phi_{s1},\Phi_{s2},...,\Phi_{st},....\Phi_{sT})$, where T is the number of stages in the planning horizon. It is often convenient to refer to the set of output variables $\Omega_{\cdot}$ that contains calculated values of input factors, production functions, attributes etc. The distinction between the elements of $\Phi_{s\cdot}$ and $\Omega_{\cdot}$ is not clear, usually $\Omega_{\cdot}$ is a subset of $\Phi_{s\cdot}$. The elements in $\Omega_{\cdot}$ will usually be traits that at least in principle can be observed in the real system. The term *model input* usually refers to (Φ_0,Θ).

The purpose of the models are to calculate the expected utility, i.e.,

$$\bar{U}(\Theta)=\int_{-\infty}^{\infty}U(\Theta,\Phi=\phi)p(\Phi=\phi)d\phi=\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}U(\Theta,\phi_{s\cdot}|\Phi_0=\phi_0)\ p(\Phi_{s\cdot}=\phi_{s\cdot}|\Phi_{\mathbf{0}}=\phi_{\mathbf{0}})\ p(\Phi_{\mathbf{0}}=\phi_{\mathbf{0}})\ d\phi_{s\cdot}\ d\phi_{\mathbf{0}} \quad \textbf{(38)}$$

where $\bar{U}(\Theta)$ is the utility function. In general it can refer to any function of the output variables.

Simulation models are numerical methods for solving this integral.

Two different categories of simulation models have been implemented within animal production. Stochastic models and deterministic models, where the stochastic nature of the system is ignored (i.e. the underlying assumption is that $P(\Phi=\phi_e)=1$ for some ϕ_e e.g. estimated from various experiments, or that $U(\Theta,\Phi)$ is linear in the parameters). Examples of deterministic models can be found in Whittemore and Fawcett (1976; Black (1995); Arnold

and Bennett (1991); Danfær (1990). Stochastic models can be further subclassified into Probabilistic Models and Monte Carlo models. Probabilistic models are models such as Markov Chain models (see references under Dynamic Programming and in addition e.g. Jalvingh *et al.*, 1993) and Bayesian Networks. Within the probabilistic models the distribution of the output variables can be directly found within a single run of the model. Reasonable complex models can be specified within this context, at least if the parameters and the traits follow the Gaussian (normal) distribution. Capacity restrictions, interactions between system elements and the inclusion of decision variables will, however, make it impossible to specify the distribution in closed form. Therefore, the Monte Carlo simulation technique is chosen. The Monte Carlo technique relies on the drawing of random numbers. Every time the model encounters a stochastic variable, a (pseudo)-random variable is drawn from the appropriate distribution and this value used in the subsequent calculations. Each completed calculation (simulation run) with the model represents a random drawing from the simultaneous distribution of input and output variables. By increasing the number of calculations the distribution of the output variables can be specified to any degree of precision. The expected utility is found from:

$$\bar{U}(\Theta)) \approx \frac{1}{k}\sum_{i=1}^{k} U(\Theta,\phi_i)$$

where ϕ_i is a random drawing from the multidimensional distribution of the parameters, and k is the number of random drawings (simulation runs). In addition the standard error on the estimated utility can be found by calculating the variance of $U(\Theta,\phi_i)$. Thus, we can obtain a measure of the precision of the estimated utility, and an indication of how many iterations that are necessary. E.g. if standard error of the expected utility is 10% after 100 iterations, it will take 10,000 iterations to obtain a standard error of 1%.

Examples of Monte Carlo simulation models are such models as Singh (1986); de Roo, (1987); Sørensen *et al.* (1992).

Simulation models can also be divided between physiological models of single animals, physiological models of whole herds, and models of whole herds with emphasis on managements strategies. The physiological model of whole herds is e.g. Tess *et al.* (1983); Pettigrew *et al.* (1986) and Finlayson *et al.* (1995), while examples of current physiological models is found under the deterministic models, and the whole herd approach under the stochastic models mentioned above. Obviously, it is within the last category that the likely candidates for decision support systems should be found. However, models from the first category have been adapted to serve as decision support systems. The first two approaches are often based on a description of the system with differential (or difference) equations, while the third approach relies more on the theory behind stochastic processes such as queuing models.

Compared to the other techniques, simulation models are much more flexible, and there is no constraint on the degree of detail in the model. Especially when the so-called object oriented programming method is used, it is possible to achieve a very close correspondence between the elements of the real system and the model. See e.g. Chang *et al.* (1994), Skidmore and Beverly, (1995); and Jørgensen and Kristensen, (1995). Any model variable can be used as output variable and it is easy to represent capacity restrictions.

Very often the purpose of simulation models is to improve the understanding of a system, i.e. to combine research results from different areas to obtain a comprehensive description of the system, the so-called holistic approach. This purpose should be seen as something different from decision support. When simulation models are used to improve the understanding of the complex system, a fixed and known set of parameters are used for the initial state of nature, Φ_0, and the expected value of the utility function or any other output variable is calculated as:

$$\bar{U}(\Theta|\Phi_o=\phi_{oj}))\approx\frac{1}{N}\sum_{i=1}^{N}U(\Theta,\phi_{si}|\Phi_o=\phi_{oj}) \qquad \textbf{(40)}$$

i.e. only the inner part of the integrand in (38) is solved.

The knowledge of the systems sensitivity to changes in the parameters is part of, what we call understanding of a system.

In contrast, when simulation models are used to determine 'optimal' strategies we want to find the optimal set of decision rules given the precision in our current knowledge of the parameters (state of nature). The parameters used in each simulation run should therefore be a sample from the apriori distribution of Φ_0 reflecting the precision in our current knowledge, and not fixed values.

The search for optimal strategies is included in linear programming, dynamic programming and influence diagrams, (i.e. simplex algorithm, policy and value iteration). No such search facility is included in connection with simulation models. This is a major drawback of the method.

Within simulation models the search for the optimal set of decision rules is treated as a general problem of multidimensional optimization. Several numerical methods exist that can handle this (see e.g. Press *et al.*, 1989). The choice of method should be made carefully. The flexible form of the simulation models means that the behaviour of the expected utility function is unknown, for example if there exist discontinuities and local optima. Such phenomenons can make some of the methods go wrong. Another complication is that the expected utility is only estimated with a precision depending on the number of simulation runs within each treatment. The difference between to set of decision rules may therefore be just a matter of sampling error, rather than a difference in expected utility. The solution to this problem is to do more simulation runs. But there is a trade off between time spent improving the precision in the estimate of one set of decisions rules and the time spent searching for a better. Guide lines to handle this problem is not available.

The search procedures are most easily demonstrated by borrowing terms from experimental world. A set of decision rule is termed a treatment. Expected values from a given treatment are found by a number of replicates (N) or simulation runs. When searching for optimal decision rules, we have to repeatedly specify new treatments and calculate expected utility for the treatment. If we want to combine a set of treatments simultaneously, we design an experiment with the different treatments included.

A well-established technique for well-behaved expected utility functions, especially with continuous variables in the decision rules, is the gradient search technique. First an experiment

is designed to initially explore the expected utility function, e.g. a response surface design. The result from this experiment is analysed and the response surface estimated. If the optimum is outside the current design, the path of steepest ascent of the response surface is estimated. Then an experiment is made with treatments on the steepest ascent path, until the optimal treatment on this path is found. A new response surface design is made centred around this optimum point. This procedure is repeated until the optimum is found with sufficient precision. Using this procedure, an (at least local) optimum will be found.

Other promising techniques are stochastic search techniques, such as *simulated annealing* and *genetic algorithms*. These algorithms start with the selecting of a random initial set of decision rules (treatment) as the current. The expected utility of this is calculated. Then the following steps are carried out iteratively. Select a new treatment candidate based on the current treatment by random permutation. Calculate expected utility for the treatment candidate. Decide randomly according to a specific rule (depending on the technique), whether to use the treatment candidate as current candidate by drawing a random variable. Continue the iterations.

A third possibility is the group of so-called heuristic search strategies. Examples can be found in Reeves (1993).

If we compare the simulation method with the information needs of the decision problem, all the aspects can be covered, and the utility function and capacity restrictions can easily be handled. The curse of dimensionality is not felt immediately. The computation time of a single run of the model grows more or less linearly with the complexity of the model. The problem is the search for optimal solutions. The techniques mentioned are not as efficient as either the simplex, value iteration or policy iteration methods. With the same complexity in the decision rules as in e.g. dynamic programming the curse of dimensionality will be felt, e.g. if the decision to cull an animal should include the states of all other animals in the herd. If the rules are specified more heuristically, such as cull the worst animal, the problems become tractable, but no overall optimal solution is guarantied. Other decision rules might exist with higher utility.

Probably because of this problem, published results from simulation model research usually have only very few options in the decision rule, and the decision rules are often of a very general nature. The use of simulation models for decision support is usually suggested to be of the *what-if* nature, i.e. the user of the model specifies some decision rules and the model calculates the expected output from these decision rules. This approach has advantages because there is no need to attempt to formulate the farmer's utility function. The user of the model can simply look at the list of output variables for different set of decision-rules and decide which set to prefer. Anyhow, it seems that some kind of optimality search within the simulation models would be the best.

The future developments within simulation modelling, will probably be in the area of estimating model parameters, perhaps by directly using the model calculations. More efficient strategies for sampling than the purely random approach, and improved search strategies are of interest, too. Finally, developments within the area of multicriteria optimisation, to obtain a better reflection of the farmers utility in the object function should not be overlooked.

References

Adler, H.C. 1978. Animal welfare in farm animal production. In: D.W. Folsch (ed.): *Ethology and Ethics of Farm Animal Production*. Basel and Stuttgart: Birkhauser.

Arnold, R.N. & G.L. Bennett. 1991. Evaluation of four simulation models of cattle growth and body Composition: Part I - Comparison and characterization of the Models. *Agricultural Systems 35*, 401-432.

Anderson, J.R., J.L. Dillon & J.B. Hardaker. 1977. *Agricultural Decision Analysis.* Ames, Iowa: The Iowa State University Press.

Bellman, R. E. (1957). *Dynamic Programming*. Princeton: Princeton University Press.

Ben-Ari, Y. and Gal, S. 1986. Optimal replacement policy for multicomponent systems: An application to a dairy herd. *European Journal of Operational Research* 23, 213-221.

Bielza, C., P. Muller, D. Insua & D.R. June. 1996. Monte Carlo Methods for Decision Analysis with Applications to Influence Diagrams.. *MCMC Preprint Service (http://www. statslab. cam. ac. uk/~mcms/html/).*

Black, J.L. 1995. The Evolution of Animal Growth Models Modelling Energy Metabolism in the Pig - Critical Evaluation of a Simple Reference Model. in *Modelling Growth in the Pig (Series: Eaap Publication: 78)*. (Moughan, P. J., Verstegen, M. W. A., and Visserreyneveld, M. I., Eds.) Wageningen Press, Stichting, Postbus 42, 6700AA, Wageningen, Netherlands. 3-9.

Brambell, F.W.R. 1965. (Chairman) *Report of the Technical Committee to enquire into the Welfare of Animals kept under Intensive Livestock Husbandry Systems.* London: HMSO, Command Paper No. 2836.

Brantas, G.C. 1975. Quoted in "Introduction to plenary session I. Intensive confinement farming: Ethological, veterinary and ethical issues". *Proceedings of the 1st World Congress of Ethology and Applied Zootech*. Madrid, 1978.

Chang, W., D.H. Streeter, & L.R. Jones. 1994. An Object-Oriented Model for Simulating Milking Parlor Operations. *Journal of Dairy Science 77*, 84-93.

Charnes, J.M. & P.P. Shenoy. 1996. A Forward Monte Carlo Method for Solving Influence Diagrams Using Local Computation (abstract). in *Proceedings of the First European Conference on Highly Structured Stochastic Systems. 19.-24. may 1996. Rebild Denmark*.

Danfær, A. 1990. A Dynamic Model of Nutrient Digestion and Metabolism in Lactating Dairy Cows. *Beretning fra Statens Husdyrbrugsforsøgs 671*.

Dawkins, M. 1980. *Animal Suffering*. Chapman and Hall: London.

de Roo, G. 1987. A stochastic model to study breeding schemes in a small pig population. *Agricultural Systems 25*, 1-25.

Dindorph, U. 1992. Agricultural applications of knowledge based system concepts - exemplified by a prototype on weed control in organic farming. WEEDOF. Ph.D. dissertation. *Tidsskrift for planteavls Sprecialserie* Beretning nr S 2201, pp. 1-97.

Dutch National Council for Agricultural Research. 1977. *Report of Committee of Experts on Animal Husbandry and Welfare: Policy Recommendations*. The Hague.

Emmans, G.C. 1995. Problems in modelling the growth of poultry. *World Poultry Science Journal 51*, 77-89.

Finlayson, J.D., O.J. Cacho, & A.C. Bywater. 1995. A simulation model of grazing sheep: I. Animal Growth and Intake. *Agricultural Systems 48*, 1-25.

Fraser, A.F. & D.M. Broom. 1990. *Farm Animal Behaviour and Welfare, 3rd edition.* London: Baillière.

Hansen, J.P. 1992. ADAM-H - et beslutningsstøttesystem til fastlæggelse af optimal foderforsyning og arealanvendelse i kvægbedriften. *Beretning fra Statens Husdyrbrugsforsøg 718*.

Hardie, A.J. 1996. Ph.D. thesis. University of Exeter.

Hayes-Roth, F., D.A. Waterman & D.B. Lenat. 1996. *Building Expert Systems*, Addison-Wesley.

Houben, E.H.P., R.B.M. Huirne, A.A. Dijkhuizen and A.R. Kristensen. 1994. Optimal replacement of mastitis cows determined by a hierarchic Markov process. *Journal of dairy Science 77*, 2975-2993.

Houben, E.H.P., R.B.M. Huirne, A.A. Dijkhuizen & J.F. Schreinemakers.1995. Genetic algorithm to link decisions at animal and herd level on dairy farms. In: Houben, E.H.P. 1995. Economic optimization of decisions with respect to dairy cow health management. I Ph.D. thesis. Department of Farm Management, Wageningen Agricultural University. pp. 101-121.

Howard, R. A. 1960. *Dynamic programming and Markov processes*. Cambridge, Massachusetts: The M.I.T. Press.

Howard, R.A. & J.E. Matheson. 1981. Influence Diagrams. In: *Readings on the Principles and Applications of Decision Analysis*. (Howard, R. A. & J.E. Matheson, Eds.) Strategic Decisions Group, Menlo Park, California. 719-762.

Hughes, B.O. 1976. Behaviour as an index of welfare. *Proceedings 5th European Conference*, Malta.

Jalvingh, A.W., J.A.M. van Arendonk & A.A. Dijkhuizen. 1993. Dynamic Probabilistic Simulation of Dairy Herd Management Practices. 1. Model Description and Outcome of Different Seasonal Calving Patterns. *Livestock Production Science 37*, 107-131.

Jensen, A.L. 1995. A probabilistic model based decision support system for mildew management in winter wheat. Ph.D. thesis. Dina Research Report No. 39. URL: ftp://ftp.dina.kvl.dk/pub/Dina-reports/research39

Jensen, F., F.V. Jensen & S.L. Dittmer. 1994. From Influence Diagrams to Junction Trees. in *Proceedings of the 10th Conference on Uncertainty in Artificial Intelligence. San Francisco, July 1994. Morgan Kaufmann.* (de Mantaras, R. L. and Poole, D. , Eds.) 367-373.

Jensen, K.K. & J.T. Sørensen. 1996. The idea of ethical accounting for a livestock farm. *Submitted for Journal Publication*.

Jørgensen, E. & A.R. Kristensen. 1995. An object oriented simulation model of a pig herd with emphasis on information flow. In: *FACTs 95 March 7, 8, 9, 1995, Orlando Florida, Farm Animal Computer Technologies Conference*. pp 206-215.

Kennedy, J. O. S. 1986. *Dynamic programming. Applications to agriculture and natural resources.* London and New York: Elsevier applied science publishers.

Kilgour, R. 1987. The application of animal behaviour and the humane care of farm animals. *Journal of Animal Science 46*, 1478-1486.

Kilgour, R. & C. Dalton. 1984. *Livestock behaviour. A practical guide.* London: Granada.

Kjærulff, U. 1995. dHugin: A computational system for dynamic time-sliced Bayesian Networks. *International Journal of Forecasting*. 1-29.

Kristensen, A.R. 1987. Optimal replacement and ranking of dairy cows determined by a hierarchic Markov process. *Livestock Production Science 16*, 131-144.

Kristensen, A. R. 1989. Optimal replacement and ranking of dairy cows under milk quotas. *Acta Agriculturæ Scandinavica* 39, 311-318.

Kristensen, A.R. 1991. Maximization of net revenue per unit of physical output in Markov decision processes. *European Review of Agricultural Economics 18*, 231-244.

Kristensen, A. R. 1992. Optimal replacement in the dairy herd: A multi-component system. *Agricultural Systems 39*, 1-24.

Kristensen, A.R. 1993. A feed ration formulation program for PC to be used in education at university level. In: E. Annevelink et al.: *Proceedings XXV CIOSTA-CIGR V Congress. Farm Planning, Labour and labour conditions, Computers in farm management.* pp 332-339.

Kristensen, A.R. & I. Thysen. 1991. Economic value of culling information in the presence and absence of a milk quota. *Acta Agriculturæ Scandinavica 41*, 129-135.

Lauritzen, S. L. 1995. Probabilistic Expert Systems. *Aalborg University. Insititute for Electronic Systems. Department of Mathematics and Computer Science. Technical Report* R-95-2013, pp. 20.

Lauritzen, S. L. & D.J. Spiegelhalter. 1988. Local computations with probabilities on graphical structures and their application to expert system (with discussion). *Journal of the Royal Statististical Society B 50*, 157-224.

Nilsson, D. 1996. An efficient algorithm for finding the M most probable configurations in Bayesian networks. *Aalborg University. Insititute for Electronic Systems. Department of Mathematics and Computer Science. Technical Report* R-96-2020, pp. 30.

Pearl, J. 1988. *Probabilistic reasoning in intelligent systems*, Morgan Kaufmann Publishers, San Mateo.

Pedersen, C. 1996. *Kan det betale sig at analysere grovfoder. Opgave i faget Videregående produktionsstyring.*, Institut for husdyrbrug of Husdyrsundhed, KVL, København.pp. 17.

Pettigrew, J.E., S.G. Cornelius, V.R. Eidman & R.L. Moser. 1986. Integration of factors affecting sow efficiency: A modelling approach. *Journal of Animal Science 63*, 1314-1321.

Press, W.H., B.R. Flannery, S.A. Teukolsky & W.T. Vetterling. 1989. *Numerical Recipes in Pascal. The art of scientific computing*, Cambridge University Press, Cambridge. pp. 759.

Randall, J.M. 1976. Environmental and welfare aspects of intensive livestock housing: a report on an American tour. *National Institute of Agricultural Engineering Silsoe Report* No. 19.

Reeves, C.R (ed) (1993). *Modern Heuristic techniques for combinatorial problems*, Halsted

Sandøe, P. & F. Hurnik (eds). 1996. Welfare of domestic animals. Concepts, theories and methods of measurement. *Acta Agriculturæ Scandinavica, Section A, Animal Science. Supplementum 27.*

Sandøe, P. & H.B. Simonsen. 1992. Assessing animal welfare: where does science end and philosophy begin? *Animal Welfare 1*, 257-267.

Shachter, R.D. 1986. Evaluating Influence Diagrams. *Operations Research 34*, 871-882.

Shenoy, P.P. 1992. Valuation-based systems for Bayesian decision analysis. *Operations Research 40*, 463-484.

Singh, D. 1986. Simulation of swine herd population dynamics. *Agricultural Systems 22*, 157-183.

Skidmore, A.L. & R.W. Beverly. 1995. Utilizing Object-Oriented Analysis, Design, and Programming Techniques to Develop a Dairy Herd Simulation Program. in *FACTs 95 March 7, 8, 9, 1995, Orlando Florida, Farm Animal Computer Technologies Conference*. pp. 190-195.

Sørensen, J.T. & C. Enevoldsen. 1991. Kvægsundhed, økonomi og mælkekvote. *Dansk Veterinærtidsskrift 74*, 298-302.

Sørensen, J.T., E.S. Kristensen & I. Thysen. 1992. A stochastic model simulating the dairy herd on a PC. *Agricultural Systems 39*, 177-200.

Tatman, J. A. & R.D. Shachter. 1990. Dynamic Programming and Influence Diagrams. *IEEE Transactions on Systems, Man, and Cybernetics 20*, 365-379.

Tess, M.W., G.L. Bennett, & G.E. Dickerson. 1983. Simulation of Genetic changes in life cycle efficiency of pork production. I. A Bioeconomic model. *Journal of Animal Science 56*, 336-353.

Whittemore, C.T. & R.H. Fawcett. 1976. Theoretical Aspects of a flexible model to simulate protein and lipid growth in pigs. *Animal Production 22*, 87.