Genet. Res., Camb. (1999), **74**, *pp.* 265–270. Printed in the United Kingdom © 1999 Cambridge University Press

Artificial selection on phenotypically plastic traits

MARK KIRKPATRICK* AND THOMAS BATAILLON
[1] *Section of Integrative Biology C0930, University of Texas, Austin, TX 78712 USA*
[2] *INRA, Domaine de Melgueil, 34130 Mauguio, France*

(*Received 29 March 1999 and in revised form 1 July 1999*)

Summary

Many phenotypes respond physiologically or developmentally to continuously distributed environmental variables such as temperature and nutritional quality. Information about phenotypic plasticity can be used to improve the efficiency of artificial selection. Here we show that the quantitative genetic theory for 'infinite-dimensional' traits such as reaction norms provides a natural framework to accomplish this goal. It is expected to improve selection responses by making more efficient use of information about environmental effects than do conventional methods. The approach is illustrated by deriving an index for mass selection of a phenotypically plastic trait. We suggest that the same approach could be extended directly to more general and efficient breeding schemes, such as those based on general best linear unbiased prediction. Methods for estimating genetic covariance functions are reviewed.

1. Introduction

Artificial selection programmes are predicated on statistical models in which the phenotypic expression of a trait is statistically decomposed into genetic and environmental components. The environmental component is typically viewed as uncontrolled noise that results from various developmental and physiological accidents. In some cases, however, a substantial fraction of this variation is caused by environmental variables that can be identified, such as food quality (for animals) or soil nitrogen (for crops). The response of a character to environmental factors is known as phenotypic plasticity.

Phenotypic plasticity raises several issues of practical importance in animal and plant breeding. How can we improve the efficiency of artificial selection using information about the species' response to environmental variation? How can we best select on a trait when individuals develop in different environments, and so vary phenotypically because of plasticity as well as genetic differences?

A key to this problem was proposed by Falconer (1952). His insight was to view the expression of a trait in two different environments as two different traits. One can then measure the genetic and phenotypic variation in each environment, and the correlations between them. Cast in this way, the question is translated into a standard problem involving selection on two correlated traits. This idea became the basis for discussions of artificial selection on traits expressed in two environments (Falconer, 1960), and was extended to allow for three or more discrete types of environments (Robertson, 1959; Via, 1987).

Many kinds of environmental variables, however, do not fall into discrete states; temperature and soil moisture content are two obvious examples. Here a genotype's response to the environmental variable is described by a function known as the reaction norm. Using Falconer's approach, the phenotype expressed in each environment can be viewed as a separate trait. When there is a continuum of environments, each genotype is capable in principle of expressing an infinite number of phenotypes. For this reason, reaction norms can be thought of as 'infinite-dimensional' traits (Kirkpatrick & Heckman, 1989; Gomulkiewicz & Kirkpatrick, 1992; Kirkpatrick & Lofsvold, 1992).

The question of how a continuous reaction norm evolves under natural selection has been studied by

* Corresponding author. Tel: +1 (512) 471 5996. Fax: +1 (512) 471-3878. kirkp@mail.utexas.edu

numerous evolutionary biologists (reviewed in Via *et al.*, 1995). The topic was studied using the infinite-dimensional approach by Gomulkiewicz & Kirkpatrick (1992) and Gilchrist (1996). The information inherent in continuous reaction norms seems not, however, to have been systematically exploited in artificial selection programmes. This paper outlines how one can design a selection index for traits that show phenotypic plasticity using the infinite-dimensional framework. We begin by reviewing the basic concepts of this framework. We then illustrate the approach by constructing a selection index for practising mass artificial selection on a phenotypically plastic trait, and suggest that the approach could be applied directly to breeding schemes based on more sophisticated designs.

2. The infinite-dimensional framework

In standard quantitative genetics, the phenotypic distribution for a set of traits is described by a mean vector, $\bar{\mathbf{z}}$, and a covariance matrix, **P**. Additive genetic variation is described by a corresponding genetic covariance matrix, **G**. Infinite-dimensional traits such as reaction norms have natural analogues for these quantities. Consider how a trait (such as metabolic rate) is expressed as a function of a continuous environmental variable (such as temperature). The *mean function*, $\bar{z}(x)$, simply gives the average value of the trait in environment x. The value of the *phenotypic covariance function*, $\mathscr{P}(x_1, x_2)$, is equal to the phenotypic covariance between the trait exposed in environments x_1 and x_2. Likewise, the value of the *additive genetic covariance function*, $\mathscr{G}(x_1, x_2)$, gives the genetic covariance between expression of the trait in those environments.

The values of some traits can be measured on an individual in each of several environments; metabolic rate as a function of temperature is an example. For other kinds of traits, however, only one measurement is possible per individual. Crop yield varies in response to soil nitrogen, for instance, but can only be measured once (at harvest time) for each individual. Conceptually, it is perhaps easiest to think about $\mathscr{G}$ and $\mathscr{P}$ in these cases as though we could estimate them directly by cloning a large number of genotypes and distributing the clones over very many environments. We discuss the practical problem of estimation below.

The descriptive statistics $\bar{z}$, $\mathscr{G}$ and $\mathscr{P}$ can be used to predict selection response if we adopt a model of inheritance. Following standard quantitative genetics (Falconer & Mackay, 1996; Lynch & Walsh, 1998), we statistically decompose the phenotype expressed in environment x into the sum of two independent parts: an additive genetic component $g(x)$ and a non-additive component $e(x)$ which is attributable to dominance and environmental factors other than the one under consideration. Both g and e are assumed to be normally (Gaussian) distributed. Then, in the absence of epistasis, the change in the mean function caused by a single generation of selection is given by:

$$\Delta\bar{z}(x) = \int\int \mathscr{G}(x, y_1)\,\mathscr{P}^{-1}(y_1, y_2)\,s(y_2)\,\mathrm{d}y_1\,\mathrm{d}y_2 \qquad (1a)$$

$$= \int \mathscr{G}(x, y)\,\beta(y)\,\mathrm{d}y, \qquad (1b)$$

where the integrals are taken over the full range of environmental states that the population experiences. In Equation (1*a*), $s()$ is the *selection differential function*, defined as the difference between the mean phenotypes that would be expressed in environment x by the selected individuals and those that would be expressed in x by all individuals. The function $\mathscr{P}^{-1}$ is the generalized inverse of the phenotypic covariance function $\mathscr{P}$ (see Kirkpatrick & Heckman, 1989).

Inspection of (1*a*) reveals that it is simply the infinite-dimensional version of the classic breeder's equation for multiple traits:

$$\Delta\bar{\mathbf{z}} = \mathbf{G}\mathbf{P}^{-1}\mathbf{s} \qquad (2)$$

(Lynch & Walsh, 1998). Loosely speaking, the summations used in matrix multiplication have been replaced by integrals. Equation (1*b*) represents the selection response in terms of the *selection gradient*, $\beta(x)$. This is a measure of the strength of directional selection (Lande & Arnold, 1983; Kirkpatrick & Heckman, 1989; Kirkpatrick, 1993; Gomulkiewicz & Beder, 1996; Beder & Gomulkiewicz, 1998). Multiplying a selection gradient by the corresponding phenotypic standard deviation gives the selection intensity, a non-dimensional quantity that is widely used by breeders (Falconer & Mackay, 1996, p. 189).

What advantages do we gain by viewing these traits as infinite-dimensional? Kirkpatrick & Heckman (1989) discuss several. Most important is that the infinite-dimensional approach increases statistical power and hence improves the accuracy of our predictions for the response to selection. Accuracy is gained because the infinite-dimensional framework makes use of the information about the ordering of the environmental states. The ordering is exploited to correct estimation errors in the genetic and phenotypic variances and covariances (see Section 4). A second advantage of the infinite-dimensional model is that it makes no *a priori* assumptions about what types of reaction norms can be produced by selection. Schemes which begin by assuming that reaction norms can be described by a particular family of functions automatically restrict the range of possible outcomes. Third, the infinite-dimensional framework naturally generates a description for the value of the trait expressed in any environment, not just in a finite set of environments. Together, these advantages can be expected to produce more accurate predictions of the response to selection. Analysis of simulated data sets

https://doi.org/10.1017/S0016672399004115 Published online by Cambridge University Press

does indeed show that the infinite dimensional method has increased statistical power over conventional matrix-based methods that discard information about the ordering of the environmental states (Kirkpatrick & Heckman, 1989).

3. Selection indices for phenotypically plastic traits

Economically important traits such as crop yield and growth rate often show phenotypic plasticity. Our question is how to select on such traits in order to maximize the rate of economic improvement. Two factors complicate the answer. First, the environment varies in space and time, often in unpredictable ways. To optimize a selection programme we must account for that variation. Secondly, each individual or genotype is typically measured in only one or a few environments, not the full range that is experienced by a population. How do we compare measurements from individuals raised in different environments when deciding whom to breed?

A genotype's response to a continuous environmental variable is described by its reaction norm. The response of a reaction norm to selection can be described and analysed using the infinite-dimensional framework (Gomulkiewicz & Kirkpatrick, 1992). Imagine that we clone an individual and allow the clones to develop in a set of environments that differ with respect to the environmental variable. Its reaction norm is then the curve that relates the value of the environment e to the average phenotype $\bar{z}(e)$ that results in that environment. To simplify notation, we assume that the trait is measured in terms of its economic value.

To maximize economic improvement, we maximize the expected rate of increase in the trait's value, averaging over the frequency with which different environments are experienced. Let $f(x)$ be the probability density that the individuals experience environment x. Denoting the mean phenotype across all environments as $\bar{Z}$, the expected global rate of economic improvement is

$$\Delta\bar{Z} = \int f(x)\,\Delta\bar{z}(x)\,\mathrm{d}x, \tag{3}$$

where $\Delta\bar{z}(e)$ is the change across one generation in the average phenotypes produced in environment e. The environmental distribution $f(\,)$ plays the same role here as the vector of economic weights that appear in selection indices for multiple traits (Falconer & Mackay, 1996). The problem is how to select so as to maximize $\Delta\bar{Z}$.

An analogous question arises in standard quantitative genetics when m traits are measured on each individual. The question then is how to combine these measurements such that selection based on this index will maximize the rate of improvement in economic value. Economic value, in turn, depends on n traits that may or may not be among those that have been directly measured. The value of the selection index for an individual is calculated as

$$I = \mathbf{b}^{\mathrm{T}}\mathbf{z}, \tag{4}$$

where $\mathbf{z}$ is the vector of n measurements for that individual and $\mathbf{b}$ is a vector of weighting factors that applies to the entire population. The rate of economic improvement is maximized with the weights

$$\mathbf{b}^{\mathrm{T}} = \mathbf{v}^{\mathrm{T}}\mathbf{G}\mathbf{P}^{-1} \tag{5}$$

(see Falconer & Mackay 1996, eqn 19.15). Here $\mathbf{P}$ is the $m \times m$ phenotypic covariance matrix, $\mathbf{G}$ is the $n \times m$ matrix whose ijth element is the additive genetic covariance between the ith measured trait and the jth trait of economic importance, and $\mathbf{v}$ is the vector whose ith element represents the economic value of trait i.

This argument can be extended directly to the case where economic value depends on a phenotypically plastic trait. In effect, we are selecting on an infinite-dimensional reaction norm that relates the economic value of a genotype to the environment in which it develops. Consider the situation where individual i is measured in set of m_{i} different environments. This number may be as small as one – for example when we know only the final yield of a crop and the average soil moisture during the growing season. In other cases, it may be possible to measure individuals in more than one environment – for example when data are available for an animal's daily growth rate as a function of its protein intake. We assume that the environmental states are known. We allow, however, for the possibility that different individuals are measured in different sets of environments. This will often be the case when the environmental variable depends on the climate or another uncontrolled factor.

The selection index for individual i is again calculated using (4). Now, however, the jth element of the vector of weights is given by

$$[\mathbf{b}_i]_j = \sum_{k=1}^{m_i}\left[\int f(x)\,\mathscr{G}(x, e_{ik})\,\mathrm{d}x\right][\mathbf{P}_i^{-1}]_{kj}. \tag{6}$$

Hence $\mathbf{P}_i$ is the $m_i \times m_i$ covariance matrix of phenotypes measured in the same environments that individual i was, and e_{ik} is the kth environmental state in which i was measured.

There is a clear correspondence between (5), which applies to a finite number of traits, and (6), its infinite-dimensional analogue. The summation involved in the matrix product $\mathbf{v}^{\mathrm{T}}\mathbf{G}$ of (5) has been replaced by an integral in (6). A second change is that the distribution of environmental states, $f(\,)$, seen in (6), plays the role of the vector of economic weights, $\mathbf{v}$, that appears in (5). A final difference is that there is no longer a single

https://doi.org/10.1017/S0016672399004115 Published online by Cambridge University Press

vector of weights **b** that is used for all individuals. This is because we have accommodated the possibility that different individuals have been measured in different environments, in which case the measurements are not on homologous traits.

Often each individual is measured in only a single environment. In that case, the selection index is calculated by simply multiplying each individual's phenotypic value by a weight that depends on the environment from which it came. The weight for individual i is given by (6) with m_i set equal to 1:

$$b_i = \frac{\int f(x)\,\mathscr{G}(x, e_i)\,\mathrm{d}x}{\sigma_i^2}, \tag{7}$$

where σ_i^2 ($= \mathscr{P}(e_i, e_i)$) is the phenotypic variance of all individuals in the environment where individual i developed. The value of this weighting function is greater when the trait is highly heritable in that environment (that is, σ_i^2 is small) and when selection in that environment causes large positive correlated selection response in environments that are encountered frequently.

Equations (6) and (7) are our main results. In practice, one would use them to calculate the weight(s) for each individual. These weights and the phenotypic measurements are then combined using (4) to give each individual's selection index score. The scores are put in rank order, and the desired upper (or lower) fraction of the population is bred to produce the next generation. This procedure is simple mass selection implemented so as to make best use of information regarding plasticity.

4. Estimating the genetic covariance function for a plastic trait

The selection index (Equations 6 and 7) involves the genetic covariance function $\mathscr{G}$, which we assumed was known in the previous example. How can one estimate $\mathscr{G}$? In this section we outline several alternative approaches that have recently been proposed.

Estimates of genetic covariances include sampling errors that are often large. For this reason, one generally prefers an estimate for the covariance function that smooths the data to some degree (Kirkpatrick *et al.*, 1990, 1994; Pletcher & Guyer, 1999). This is the point at which the infinite-dimensional framework makes use of the ordering of the environmental states. The sampling error in a covariance estimate is decreased by the smoothing, which makes use of the estimates from covariances corresponding to neighbouring environmental states. That is not true of traditional methods that represent the phenotype measured in each of several environments as a vector whose elements have no natural relationship to each other. Under traditional methods, one obtains the same results if the ordering of elements of the vector (corresponding to the different environmental states) is randomized. Various goodness-of-fit tests, for example the likelihood ratio test (Meyer & Hill, 1997), can be used to determine the appropriate degree of smoothing.

Smoothing requires that we choose a mathematical representation of the covariance function. The most flexible approach is non-parametric in spirit: $\mathscr{G}$ is not assumed to be of any particular form *a priori*. One approach here is to use polynomials to represent the covariance function, which which guarantees that any form of $\mathscr{G}$ is possible (Kirkpatrick *et al.*, 1990). The entries of an estimated covariance matrix **G** can be interpolated to arrive at an estimate for the covariance function $\mathscr{G}$ with an appropriate amount of smoothing (Kirkpatrick *et al.*, 1990, 1994). A mathematically equivalent approach is to use the method of 'random regression' with polynomials as the basis (Jones *et al.*, 1999; Hill & Brotherstone, 1999). Meyer & Hill (1997) developed a direct method based on restricted maximum likelihood (REML). They bypass the intermediary of a covariance matrix and pass directly from the phenotypic measurements to an estimate of $\mathscr{G}$. In addition to being more elegant, their method offers increases in statistical power. For example, when estimating the covariance matrix **G**, several similar environments might be lumped into a single environmental state to simplify the analysis. The resulting artefacts are avoided by Meyer & Hill's direct method.

There are two prices to be paid for the flexibility of a non-parametric approach. First, the resulting covariance function is typically quite 'wiggly' – a behaviour that is generic to higher-order polynomials. That problem can be rectified by using a different method of interpolation. The method of smoothing by splines (see, for example, Press *et al.*, 1992) is a very promising candidate here (Gomulkiewicz & Kirkpatrick, 1992; White *et al.*, 1998; Hill & Brotherstone, 1999). A second difficulty with the non-parametric approach is that a large number of parameters must be estimated. This leads to large errors in the estimates and produces serious computational challenges.

These problems motivate two alternative parametric approaches to the estimation of $\mathscr{G}$. The first parametric approach assumes that the covariance function can be described by a simple mathematical function (Pletcher & Guyer, 1999). For example, one might assume that the genetic correlation between any two environmental states depends only on the difference between them, and not on their absolute values. One could further suppose that the correlation declines as this difference increases according to a exponential function, say, or a Gaussian curve. In that case, we need only estimate the genetic variance in each environment and a single

https://doi.org/10.1017/S0016672399004115 Published online by Cambridge University Press

parameter describing the rate at which the genetic correlation falls off between different environments. The number of parameters to be estimated is therefore greatly reduced compared with the non-parametric approach. A further advantage is that an appropriate choice for the functional form will guarantee that the estimate of $\mathscr{G}$ is positive semi-definite, as required by the definition of a covariance function. That result is not automatic for the non-parametric methods. The drawback of this parametric approach is that the choice of functional form is arbitrary.

A second parametric approach involves smoothing the data at the level of the individual, before the genetic covariance function $\mathscr{G}$ is estimated. If individuals can be measured in each of several environments, their reaction norms can be estimated using a random regression model (see Schaeffer & Dekkers (1994) and Jones *et al*. (1999)). The model consists of a linear combination of basis functions which themselves can take any form. By choosing a sufficiently simple model (that is, one where the reaction norm is described by a weighted sum of a small number of basis functions), the number of parameters that must be estimated will be small. A major drawback of the non-parametric approach is therefore overcome. This second parametric method does, however, have weaknesses. The choice of basis functions is *ad hoc*. It is not clear how to decide the optimal degree of smoothing at the level of the individual, where often we may have only a single measurement per environment. Last, the method can only be used when individuals can be measured in multiple environments. It could be applied to daily animal growth rate as a function of nutritional quality, for example, but not to crop yield as a function of average soil moisture.

The question of how to estimate a genetic covariance function is difficult. While several approaches have been proposed, each has its drawbacks. There is no clear consensus at present about when each approach will be superior, or what the relative efficiencies of the approaches are in different situations. Further advances in this area are to be expected in the near future.

5. Discussion

The 'environmental' component of phenotypic variation has itself a genetic component. Information on the genetics of environmental sensitivity can be incorporated into standard statistical methods for breeding in order to improve the efficiency of artificial selection. Falconer's insight that the expression of a trait in a different environment can be viewed as a different trait provides a natural framework for this extension. The results above show that this basic idea can be used when the environmental state varies in a continuous fashion. These ideas could easily be extended to the case of multiple traits, or to the case of a plastic trait whose expression changes with age.

We have assumed that the frequency of different environmental states acts in effect as an economic weight. The expected economic gain that is maximized in the model above is simply the average value of the trait expected in each environment, averaged over all possible environmental states. In practice, however, it might be necessary to transform either the trait value or the environmental frequency distribution to arrive at the economic value.

The selection index developed here is a simple illustration of how information about phenotypic plasticity can be exploited in artificial selection. Modern selection programmes use more powerful designs than mass selection. A more sophisticated approach would be to expand the statistical model to account for additional fixed effects, such as the year and farm in which each measurement is taken. The widely used best linear unbiased prediction (BLUP) method, often implemented in animal breeding using the 'animal model' (Nicholas, 1987; Falconer & Mackay, 1996, p. 244), would be an efficient framework in which to do that. Ultimately, one wants to incorporate all the environmental and phenotypic data into a single statistical model in order to make breeding decisions. REML is an attractive approach here (Falconer & Mackay, 1996, p. 244). Meyer & Hill (1997) pioneered the application of REML to infinite-dimensional traits by using polynomials to find a non-parametric estimate of a covariance function. Pletcher & Guyer (1999) used likelihood to find a parametric estimate of a covariance function. An important task for the future is to expand the repertoire of REML methods to include all the approaches for estimating covariance functions that have recently been developed.

This paper is dedicated to Douglas Falconer, whose research and writings on quantitative genetics have inspired many breeders and evolutionary biologists. We thank Bill Hill, Toby Johnson and two anonymous reviewers for helpful comments on the manuscript. We are grateful for support from the National Science Foundation.

References

Beder, J. H. & Gomulkiewicz, R. (1998). Computing the selection gradient and evolutionary response of an infinite-dimensional trait. *Journal of Mathematical Biology* **36**, 299–319.

Falconer, D. S. (1952). The problem of environment and selection. *American Naturalist* **86**, 293–298.

Falconer, D. S. (1960). Selection of mice for growth on high and low planes of nutrition. *Genetical Research* **1**, 91–113.

Falconer, D. S. & Mackay, T. F. C. (1996). *Introduction to Quantitative Genetics*, 4th edn. Harlow, Essex: Longman.

Gilchrist, G. W. (1996). A quantitative genetic analysis of thermal sensitivity in the locomotor performance curve of *Aphidius ervi*. *Evolution* **50**, 1560–1572.

Gomulkiewicz, R. & Beder, J. H. (1996). The selection gradient of an infinite-dimensional trait. *SIAM Journal of Applied Mathematics* **56**, 509–523.

Gomulkiewicz, R. & Kirkpatrick, M. (1992). Quantitative genetics and the evolution of reaction norms. *Evolution* **46**, 390–411.

Hill, W. G. & Brotherstone, S. (1999). Advances in methodology for utilizing sequential records. In *Metabolic Stress in Dairy Cows*. BSAS Occasional Publication, No. 24. Ed. J. D. Oldham, G. Simm, A. F. Groen, B. L. Nielsen, J. E. Pryce and T. L. J. Lawrence, pp. 55–61.

Jones, H. E., White, I. M. S. & Brotherstone, S. (1999). Genetic evaluation of Holstein Friesian sires for daughter condition-score changes using a random regression model. *Animal Science* **68**, 467–475.

Kirkpatrick, M. (1993). Evolution of size and growth in fisheries and other harvested natural populations. In *The Exploitation of Evolving Resources* (ed. K. Stokes, J. M. McGlade & Law, R., pp. 145–154. Springer Lecture Notes in Biomathematics 99. Berlin: Springer.

Kirkpatrick, M. & Heckman, N. (1989). A quantitative-genetic model for growth, shape, reaction norms, and other infinite-dimensional characters. *Journal of Mathematical Biology* **27**, 429–450.

Kirkpatrick, M. & Lofsvold, D. (1992). Measuring selection and constraint in the evolution of growth. *Evolution* **46**, 954–971.

Kirkpatrick, M., Lofsvold, D. & Bulmer, M. (1990). Analysis of the inheritance, selection, and evolution of growth trajectories. *Genetics* **124**, 979–993.

Kirkpatrick, M., Hill, W. G. & Thompson, R. (1994). Estimating the covariance structure of traits during growth and aging, illustrated with lactation in dairy cattle. *Genetical Research* **64**, 57–69.

Lande, R. & Arnold, S. J. (1983). The measurement of selection on correlated characters. *Evolution* **37**, 1210–1226.

Lynch, M. & Walsh, B. (1998). *Genetics and Analysis of Quantitative Traits*. Sunderland: Sinauer.

Meyer, K. & Hill, W. G. (1997). Estimation of genetic and phenotypic covariance functions for longitudinal or 'repeated' records by restricted maximum likelihood. *Livestock Production Science* **47**, 185–200.

Nicholas, F. W. (1987). *Veterinary Genetics*. Oxford: Clarendon Press.

Pletcher, S. & Guyer, C. (1999). Quantitative genetics and age-dependent traits: modeling the character process. *Genetics*, in press.

Press, W. H., Teukolsky, S. A., Vetterling, W. T. & Flannery, B. P. (1992). *Numerical Recipes in C: The Art of Scientific Computing*, 2nd edn. Cambridge: Cambridge University Press.

Robertson, A. (1959). The sampling variance of the genetic correlation coefficient. *Biometrics* **15**, 469–485.

Schaeffer, L. R. & Dekkers, J. C. M . (1994). Random regressions in animal models for test-day production in dairy cattle. *Proceedings of the Fifth World Congress of Genetics Applied to Livestock Production*, *Guelph* **18** 443–446.

Via, S. (1987). Genetic constraints on the evolution of phenotypic plasticity. In *Genetic Constraints on Adaptive Evolution* (ed. V. Loeschcke), pp. 46–71. Berlin: Springer.

Via, S., Gomulkiewicz, R., de Jong, G., Scheiner, S. M., Schlichting, C. D. & van Tienderen, P. H. (1995). Adaptive phenotypic plasticity: consensus and controversy. *Trends in Ecology and Evolution* **10**, 212–217.

White, I. M. S., Thompson, R. & Brotherstone, S. (1999). Genetic and environmental smoothing of lactation curves with cubic splines. *Journal of Dairy Science* **82**, 632–638.