
Formal Darwinism, the individual-as-maximizing-agent analogy and bet-hedging

Alan Grafen

St John's College, Oxford OX1 3JP, UK (alan.grafen@sjc.ox.ac.uk)

The central argument of *The origin of species* was that mechanical processes (inheritance of features and the differential reproduction they cause) can give rise to the appearance of design. The 'mechanical processes' are now mathematically represented by the dynamic systems of population genetics, and the appearance of design by optimization and game theory in which the individual plays the part of the maximizing agent. Establishing a precise individual-as-maximizing-agent (IMA) analogy for a population genetics system justifies optimization approaches, and so provides a modern formal representation of the core of Darwinism. It is a hitherto unnoticed implication of recent population-genetics models that, contrary to a decades-long consensus, an IMA analogy can be found in models with stochastic environments (subject to a convexity assumption), in which individuals maximize expected reproductive value. The key is that the total reproductive value of a species must be considered as constant, so therefore reproductive value should always be calculated in relative terms. This result removes a major obstacle from the theoretical challenge to find a unifying framework which establishes the IMA analogy for all of Darwinian biology, including as special cases inclusive fitness, evolutionarily stable strategies, evolutionary life-history theory, age-structured models and sex ratio theory. This would provide a formal, mathematical justification of fruitful and widespread but 'intentional' terms in evolutionary biology, such as 'selfish', 'altruism' and 'conflict'.

Keywords: Darwinism; maximizing agent; game theory; bet-hedging; reproductive value; population genetics

1. INTRODUCTION

The individual-as-maximizing-agent (IMA) analogy is a pervasive and powerful tool that defines the study of adaptation as a branch of modern evolutionary biology. It is a modern, quantitative version of the problem–solution approach familiar from functional anatomy and functional morphology (Williams 1966; Rose & Lauder 1996).

It needs to be distinguished from other uses of maximization in biology. For example, part of the 'levels of selection' debate (Dawkins 1982) can be understood as whether the individual or the group or the species should be considered as the 'designed object'. Here I assume the individual. The product theorem of MacArthur (1965) is a maximization principle about a population, not an individual, as is the geometric mean principle that applies to stochastic environments, as we shall see later. Another important maximization principle in biology concerns the fundamental theorem of Fisher (1930) and the adaptive landscapes of Wright (1969). Both authors placed great significance on discovering a principle, although with differing models and differing emphases on approximate solutions, that mean population fitness was increased by natural selection. This was interpreted in two ways. First, both Fisher and Wright were keen to establish a biological analogy to free energy, which always decreases in many mechanical systems, and to entropy, which always increases in many thermodynamic systems. This kind of quantity contains important information, for example

about the positions of equilibria, and the Lyapunov function is a fundamental idea in the theory of dynamic systems (Goldstein 1980). For a modern account of biological systems, and a discussion of when the idea that gene frequencies 'hill-climb' can be strengthened to them taking the 'path of steepest ascent', see Hofbauer & Sigmund (1988). There is a second way in which Fisher (1930) and Wright (1969) interpreted the result that mean fitness increases under natural selection. Equilibrium implies no further change, which implies mean fitness cannot be increased any further. It follows that at equilibrium each individual has the highest possible mean fitness (at least locally), which can reasonably be interpreted as being at a (perhaps only local) optimum of adaptation to the environment. This is an informal version of the IMA analogy. The understandable excitement of Fisher (1930) and Wright (1969) arose from this dual role of mean fitness. Fisher (1930) employed the IMA analogy in his discussions of sex ratio, sexual selection, mimicry, and age-specific selection; used reproductive value as the maximand; and specifically discussed the theoretical conclusion that it was individual maximization and not group or species maximization that was involved. Reproductive value is a measure of the contribution made by an individual to the gene pool in distant generations.

The theoretical achievement of papers on social behaviour by Hamilton (1964) was to extend the mean fitness result, in both original roles, to cases where an

individual's genotype could affect others' numbers of offspring. The extension was inspired by the approach of Kingman (1961), and involved constructing 'inclusive fitness'. In one role this was the new function of state that was always increased during changes in gene frequencies. In the other, inclusive fitness was the new maximand—the generalization of Darwinian fitness that individuals were expected to be designed to maximize.

Once models have frequency-dependent fitnesses, however, the two roles diverge, and indeed there may be no way to define a hill that gene frequencies climb (Moran 1964). This paper is concerned with the other role, and the question of how widely the IMA analogy can be sustained. Evolutionarily stable strategy (ESS) theory is concerned with frequency dependence. Although ambivalent in applications between population genetics models, and starting straight away with a game theory model, the formal theory of ESSs takes the latter course, and simply assumes that some quantity is maximized by each individual in choosing from a defined set of options. This leaves open the fundamental questions of whether and when population-genetics systems will really result in the same equilibria as the game theory model, and how to choose the maximand so that it will be so. These same questions can also be asked of IMA analogies in inclusive fitness models, evolutionary life-history models, age-structured models, and sex ratio theory. The thesis of the present paper is that those two questions for these different areas of biology are all linked, are properly and fruitfully considered together, and represent the core of Darwinism.

2. THE IMA ANALOGY IN STOCHASTIC ENVIRONMENTS

The conceptual foundation of optimization theories has often been discussed in terms of 'approximations' and 'heuristics' and 'simplifications' (Maynard Smith 1978; Maynard Smith & Parker 1990; Seger & Stubblefield 1996). But there is fundamental theory being done, establishing formally the validity and exact nature of the analogy, which clarifies and extends Darwin's insight, the significance of which as a body of work has not been recognized. There is recent general work for inclusive fitness (Taylor 1990, 1996) extending the earlier results of Hamilton (1964, 1970), for ESSs (Hammerstein 1996; Weissing 1996), for age-structured models (Charlesworth 1994) and for frequency-dependent life histories (McNamara 1991, 1993, 1997). The IMA analogy sometimes has to be derived from population genetics for individual models where its existence or nature is likely to be controversial (Grafen 1990, 1998).

The IMA analogy is powerful, and has attracted biologists from Darwin (1859) and Fisher (1930) onwards, for a number of reasons. It abstracts from the multiplicity of loci and alleles, and encapsulates how we can expect natural selection to act on individuals. (When different parts of the genome lead to different maximands, we have intragenomic conflict.) It permits an interpretation of features of the individual organism as solutions to a problem and encourages the precise definition of that problem. The use of 'intentional language' (Dennett 1995; Rose & Lauder 1996) has proved practically indispensable

in modern evolutionary biology, where terms such as 'altruism', 'selfish', and 'conflict' define questions of interest. The dangers of anthropomorphism are more effectively avoided by linking these terms explicitly to the appropriate IMA analogy than by making valiant but doomed efforts to eliminate the use of intentional terms.

Many modelling areas have developed an IMA analogy as a valuable aid to interpretation and understanding. The most glaring exception, and therefore a major impediment to the programme of providing a general result, has been models with environmental stochasticity. Indeed, there has been a consensus for some time among biologists in different subdisciplines that such models are different, and offer no simple interpretation in terms of individual optimization (Seger & Brockmann 1987; Seger & Stubblefield 1996). In the jargon of evolutionary models, variation between (otherwise) identical individuals alive at the same time is 'demographic stochasticity', while variation across time periods affecting all individuals is 'environmental stochasticity'.

Three recent papers simultaneously presented an IMA analogy for stochastic environments (McNamara 1995; Sasaki & Ellner 1995; Haccou & Iwasa 1995), with the same basic result, but did not notice that the maximand could be interpreted as expected reproductive value.

(a) *The McNamara model*

McNamara (1995) provides (pp. 186–187) a mathematically secure treatment of the basic case, for an environment that varies between a number of environmental states s , as follows.

'The population is censused yearly. At the census point all population members are in the same given state. At this time each must choose an action from the set of possible alternatives. This choice must be made without knowledge of the environmental state in the future. The action and the environmental state during the following year determine the expected number of descendants present at the next census point in a year's time The set of pure actions available to an organism is denoted by A I define a strategy P as specifying the probability $P(a)$ that each possible pure action a will be chosen If the organism chooses action a and the ensuing environmental state is s the organism leaves $d(a,s)$ expected descendants at the census point in a year's time.'

McNamara goes on to introduce the 'profile' of a strategy P as a function $r_P(s)$ denoting the average success of P when the next environmental state is s , formally

$$r_P(s) = \int_A d(a,s) dP(a).$$

The states are treated formally as non-negative real numbers, but the analysis applies to discrete and multi-dimensional states too. Crucial to McNamara's model is the 'convexity assumption': if two possible mutations specify two phenotypes, then every probabilistic mixture of those phenotypes is also specified by its own possible mutation. It will often be reasonable to proceed on the basis of this untrue assumption.

McNamara (1995) then employs the mean of the logarithm of fitness to derive the evolutionary equilibrium following population-genetics arguments based on invasibility. Here I shall instead pursue the IMA analogy, and construct expected reproductive value, with the aim of arriving at the same conclusions by a different route.

The new point is that the total reproductive value of all individuals of a species must be considered as constant, parallel to sex ratio theory for diploids, the starting point of which is that the total reproductive value for each sex is considered to be fixed (Fisher 1930). This makes sense, because to contribute one offspring when the population contains ten individuals gains 10% of the species' reproductive value, which will have a much more positive effect on the individual's number of descendants in the long run than contributing 1000 offspring at a moment when the total population is 1 000 000 individuals, which gains only 0.1%.

The expected reproductive value therefore depends on how many offspring other individuals have in each environmental state. The total population size equals the number of parents multiplied by the average fecundity. So, taking an ESS approach (Maynard Smith & Price 1973; Maynard Smith 1982), the reproductive value in state s of an individual playing P when the rest of the (very large) population is playing P^* is proportional to $r_P(s)/r_{P^*}(s)$, where the total number of individuals in the parental generation can safely be omitted from the denominator because it is constant across states. The arithmetic average of reproductive value can then be expressed as

$$\mathcal{Z}(P, P^*) = \int_0^\infty \frac{r_P(s)}{r_{P^*}(s)} f(s) ds,$$

where $f(s)$ is the probability distribution of s .

This formula arises in McNamara's Theorem 2 (his equation 14), which shows that P^* is an optimal strategy if and only if

$$\int_0^\infty \frac{r_P(s)}{r_{P^*}(s)} f(s) ds \leq 1 \quad \text{for all strategies } P,$$

showing (because the integral equals unity when $P = P^*$) that the optimal strategy is the solution to a game with $\mathcal{Z}(P, P^*)$ as the pay-off function. This establishes the IMA analogy with expected reproductive value as maximand, for the dynamic system analysed by McNamara (1995).

The generality of McNamara's model and the simplicity of the argument suggest that the result is a general one. Extending it to more complex cases will require a careful construction of reproductive value, perhaps, for example, having spatially varying reproductive values in spatial models. For recent theoretical work on defining reproductive values in different contexts, see for example Tuljapurkar (1989), Boomsma & Grafen (1991), Taylor (1996), McNamara (1997) and Grafen (1998).

(b) *The Sasaki & Ellner and Haccou & Iwasa models*

Sasaki & Ellner (1995) note that the result shows individuals acting to maximize relative fitness, but contrast this with the absolute fitness they say is generally held to be maximized. Where in the past absolute number has been successfully used, this is because the total

number of offspring has been fixed, a special case in which the 'absolute' and 'relative' criteria are equivalent.

Haccou & Iwasa (1995) obtain the result, and go on to show that when individuals possess partial information about the environmental state, selection leads to relative fitness being maximized conditional on that information. This is one of the very attractive features to be expected from a natural IMA analogy, and not shared by the product theorem of MacArthur (1965), or the geometric mean result for stochastic environments.

3. BET-HEDGING

The above results hence show that a simple IMA analogy works in this case. Bet-hedging is therefore not needed, but is it an equally good, alternative interpretation? While the maximization by the successful genotype (i.e. the aggregate of individuals sharing that genotype, not an individual possessing that genotype) of geometric mean is a sound conclusion from population genetics, and so is not in question here, 'bet-hedging' is an intentional term that invites judgment as an IMA analogy. A simple example will now show that bet-hedging fails in the role of yielding insight into the decisions of an individual in terms of the individual's own situation.

Suppose an infinite asexual species has discrete generations, and that each generation can take place in a wet or a dry environment. Suppose the odds are equal, and generations have independent chances. An individual can allocate a fraction μ of her reproductive effort to offspring who will survive in a wet environment but die in a dry one, the remaining $1 - \mu$ going on offspring who will survive in a dry environment, but die in a wet one. If only pure strategies are allowed, then the theory above shows that the ESS strategy is equal allocation between wet and dry offspring, so the ESS strategy has $\mu = 1/2$.

Now consider a class of mixed strategies, and make the usual assumption that each randomization by an individual in implementing her mixed strategy is independent of all other randomizations. Suppose $\mathbb{I}(t)$ produces $\mu = 1/2 - t$ with probability $1/2$, and $\mu = 1/2 + t$ with probability $1/2$. The ESS in the pure case can thus be denoted $\mathbb{I}(0)$. The mean and variance of the number of surviving offspring are readily calculated to be $1/2$ and t^2 , respectively. If the variance of an individual's number of offspring played the role suggested by bet-hedging, then $\mathbb{I}(0)$ should be advantageous compared to higher values of t . In fact, these strategies have the same mean relative fitness and so, according to the result presented earlier, are selectively neutral. Any one, or any mixture, of the $\mathbb{I}(t)$ strategies will therefore be neutrally stable in the population. Further, any one of the $\mathbb{I}(t)$ strategies would invade a population the mean of which was not equal to $1/2$, regardless of the variances involved. We may conclude that the significance imputed by bet-hedging to the variability of an individual's number of offspring is actually wrong.

The variability of the total offspring number of all individuals sharing the same genotype does have significance, by the geometric mean result, but this body of individuals does not make decisions, and it is not at the level of that body of individuals that we recognize design and adaptation. It may also be worth noting that some statistical

non-independence between the decision making of individuals may arise from taking decisions conditional on an environmental feature such as whether it rains or shines on 1 May. This can be incorporated into the model by including that environmental feature in the description of the environmental state.

Population genetics models with stochastic environments are technically unaffected by these results, but they are now susceptible to reinterpretation. It has often been noted that successful genotypes are those that maximize geometric mean fitness across environmental stochasticity, which is contrasted with the maximization of arithmetic mean fitness across demographic stochasticity. The new interpretation is that the successful genotype maximizes the arithmetic mean of reproductive value in both cases. Effectively, the mathematics above show that the strategy that maximizes the geometric mean of absolute number of offspring is the same as the strategy that maximizes the arithmetic mean of the relative number of offspring (relative to itself). The first approach is natural in dynamic systems, and in population genetics; the second is natural in optimization theories and in the study of adaptation. Biologists need to appreciate both approaches and the link between them.

Seeger & Brockman (1987) present two main, theoretically well-supported, conclusions in their review of 'bet-hedging', which can now be reinterpreted in terms of individuals maximizing their expected reproductive value. The first is that organisms should not 'put all their eggs in one basket', that is, they should behave more variably than would be expected from maximizing expected fitness alone. The new interpretation is that in a variable environment, there are different states that may arise. If the population produces very few offspring in the event that a particular state arises, an opportunity may arise for a mutant to make great proportional gains by choosing to produce offspring in that state. As an equilibrium strategy is defined by being immune to invasion by all mutants, it must produce enough offspring in each possible state to ensure that none of those opportunities is worthwhile. If the environment becomes more variable in the sense that no pure strategy can ensure a reasonable number of offspring in each environmental state, then a strategy can be uninvincible only by being mixed (in a way reminiscent of the maintenance of genetic variability and sexual reproduction by environmental variability and fluctuations (Bell 1982; Hamilton *et al.* 1990)), for then, aggregating across all its players, the strategy can sufficiently populate each environmental state. A stable strategy must 'cover all bases'.

The second conclusion is that a 'bird in the hand is worth two in the bush', that is, an organism may be prepared to accept a lower mean fitness in order to reduce its variance in fitness. The appeal of this conclusion is its analogy with sound financial practice, and with the portfolio theory that underlies it (Markowitz 1991). As shown above, the avoidance of variance at the individual level is not always present. There is, however, a crude trend that can be interpreted in a new light. A population of identical individuals playing a pure strategy, in which there is high variance in offspring number between environmental states, will tend to imply there are very many offspring present in some environ-

mental states, and very few offspring present in others. States with few offspring represent exploitable opportunities for mutants, and so there will be a tendency for equilibrial, uninvincible strategies not to have states with few offspring. States with many offspring are hardly worth placing offspring in, because each is worth, proportionately, so little. Hence equilibrium strategies will tend not to have very under-populated or very over-populated states, and so will tend to have a low variance.

A further curiosity of environmentally stochastic models, and a *prima facie* difficulty for adaptationism, has been 'coin-flipping altruism' (Cooper & Kaplan 1982), the idea that within a mixed strategy some of the pure strategies are better off than others. Some offspring of a 'bet-hedging' parent are thought to be having their fitness sacrificed in order to improve the net fitness of the whole sibship. We have seen that individuals at equilibrium will maximize their expected reproductive value, in this case their relative number of offspring. A mixed strategy must therefore ensure equal expected relative offspring number for the supporting pure strategies (Sasaki & Ellner 1995). In general this will imply that the absolute number of offspring of the pure strategies will be different, and so some will seem to be worse off than others. Seeger & Brockman (1987, p. 193) write, 'This is what distinguishes genuine risk spreading from a mere polymorphism maintained by frequency-dependent selection; some of the equilibrium phenotypes may be truly worse off than others, by almost any standard'. But expected reproductive value is the natural standard, and we have seen that the equilibrium phenotypes are indeed all equally well-off when judged by it. Informally, an offspring produced for a rare and little-populated eventuality may regret that their day of glory is unlikely to arise, but this will be exactly compensated by the extent of their relative triumph should it do so.

The present paper can find no intrinsic theoretical significance for either the mean or the variance of an individual's absolute number of offspring. The mean of reproductive value, properly understood as relative to the population, is maximized at equilibrium in a wide class of models, and the variance of a strategy's reproductive value is irrelevant. It may be concluded that the concept of 'bet-hedging' is not helpful when considering individuals as decision makers.

4. CONCLUSIONS

The very success of the influential verbal arguments (Darwin 1859; Dawkins 1976; Dennett 1995) that lead many biologists to believe that 'organisms maximise their fitness' appears to have engendered resistance (Allen *et al.* 1998), and this has been unfortunate. The conceptual framework outlined here details the important theoretical challenge of producing a unifying theory to link a general population-genetics model with a general IMA game theory model. The main tasks are to find a suitably general population genetics model and a suitably general definition of reproductive value. Key elements must be the incorporation of social actions and a generalization of reproductive value to encompass inclusive fitness. Lack of technical clarity in the past explains why the IMA analogy may have been confused with other maximizing

principles, hindering theoretical development. The incorporation of a class of stochastic environment models into the simple framework of maximizing the arithmetic average of reproductive value has removed a major obstacle from the 'formal Darwinism unification project', extending our understanding of reproductive value, and reconciling previously puzzling conclusions with the main body of evolutionary theory.

I am grateful for the helpful advice of Richard Dawkins, Bill Hamilton, Olof Leimar, Mark Ridley, Steven Siller and Peter Taylor.

REFERENCES

- Allen, C., Bekoff, M. & Lauder, G. V. 1998 *Nature's purposes: analyses of function and design in biology*. Cambridge, MA: MIT Press.
- Bell, G. 1982 *The masterpiece of nature*. London: Croom Helm.
- Boomsma, J. J. & Grafen, A. 1991 Colony-level sex ratio selection in the eusocial Hymenoptera. *J. Evol. Biol.* **3**, 383–407.
- Charlesworth, B. 1994 *Evolution in age-structured populations*. Cambridge University Press.
- Cooper, W. S. & Kaplan, R. H. 1982 Adaptive 'coin-flipping': a decision-theoretic examination of natural selection for random individual variation. *J. Theor. Biol.* **94**, 135–151.
- Darwin, C. R. 1859 *The origin of species*. London: John Murray.
- Dawkins, R. 1976 *The selfish gene*. Oxford University Press.
- Dawkins, R. 1982 *The extended phenotype*. Oxford: W. H. Freeman.
- Dennett, D. C. 1995 *Darwin's dangerous idea*. London: Allen Lane, Penguin.
- Fisher, R. A. 1930 *The genetical theory of natural selection*. Oxford University Press.
- Goldstein, H. 1980 *Classical mechanics* (2nd edn). New York: Addison-Wesley.
- Grafen, A. 1990 Sexual selection unhandicapped by the Fisher process. *J. Theor. Biol.* **144**, 475–518.
- Grafen, A. 1998 Fertility and labour supply in *Femina economica*. *J. Theor. Biol.* **194**, 429–455.
- Haccou, P. & Iwasa, Y. 1995 Optimal mixed strategies in stochastic environments. *Theor. Popul. Biol.* **47**, 212–243.
- Hamilton, W. D. 1964 The genetical evolution of social behaviour. *J. Theor. Biol.* **7**, 1–52.
- Hamilton, W. D. 1970 Selfish and spiteful behaviour in an evolutionary model. *Nature* **228**, 1218–1220.
- Hamilton, W. D., Axelrod, R. & Tanese, R. 1990 Sexual reproduction as an adaptation to resist parasites (a review). *Proc. Natl Acad. Sci. USA* **87**, 3566–3573.
- Hammerstein, P. 1996 Darwinian adaptation, population-genetics and the streetcar theory of evolution. *J. Math. Biol.* **34**, 511–532.
- Hofbauer, J. & Sigmund, K. 1988 *The theory of evolution and dynamical systems*. Cambridge University Press (originally published in 1984 in German).
- Kingman, J. F. C. 1961 A mathematical problem in population genetics. *Proc. Cambridge Philos. Soc.* **57**, 574–582.
- MacArthur, R. H. 1965 Ecological consequences of natural selection. In *Theoretical and mathematical biology* (ed. T. H. Waterman & H. J. Morowitz), pp. 388–397. New York: Blaisdell.
- McNamara, J. M. 1991 Optimal life histories: a generalisation of the Perron–Frobenius theorem. *Theor. Popul. Biol.* **40**, 230–245.
- McNamara, J. M. 1993 State-dependent life-history equations. *Acta Biotheor.* **41**, 165–174.
- McNamara, J. M. 1995 Implicit frequency dependence and kin selection in fluctuating environments. *Evol. Ecol.* **9**, 185–203.
- McNamara, J. M. 1997 Optimal life histories for structured populations in fluctuating environments. *Theor. Popul. Biol.* **51**, 94–108.
- Markowitz, H. M. 1991 *Portfolio selection: efficient diversification of investments* (2nd edn). Oxford: Blackwell.
- Maynard Smith, J. 1978 Optimization theory in evolution. *Ann. Rev. Ecol. Syst.* **9**, 31–56.
- Maynard Smith, J. 1982 *Evolution and the theory of games*. Cambridge University Press.
- Maynard Smith, J. & Parker, G. A. 1990 Optimization in biology. *Nature* **348**, 27–33.
- Maynard Smith, J. & Price, G. R. 1973 The logic of animal conflict. *Nature* **246**, 15–18.
- Moran, P. A. P. 1964 On the nonexistence of adaptive topographies. *Ann. Hum. Genet.* **27**, 383–393.
- Rose, M. R. & Lauder, G. V. (eds) 1996 *Adaptation*. London: Academic Press.
- Sasaki, A. & Ellner, S. 1995 The evolutionarily stable phenotype distribution in a random environment. *Evolution*, **49**, 337–350.
- Seger, J. & Brockmann, H. J. 1987 What is bet-hedging? *Oxford Surv. Evol. Biol.* **4**, 182–211.
- Seger, J. & Stubblefield, J. W. 1996 Optimization and adaptation. In *Adaptation* (ed. M. R. Rose & G. V. Lauder), pp. 93–123. London: Academic Press.
- Taylor, P. D. 1990 Allele-frequency change in a class-structured population. *Am. Nat.* **135**, 95–106.
- Taylor, P. D. 1996 Inclusive fitness arguments in genetic models of behaviour. *J. Math. Biol.* **34**, 654–674.
- Tuljapurkar, S. 1989 An uncertain life: demography in random environments. *Theor. Popul. Biol.* **35**, 275–294.
- Weissing, F. J. 1996 Genetic versus phenotypic models of selection: can genetics be neglected in a long-term perspective? *J. Math. Biol.* **34**, 533–555.
- Williams, G. C. 1966 *Adaptation and natural selection*. Princeton University Press.
- Wright, S. 1969 *Evolution and the genetics of populations*. University of Chicago Press.

