against HIV is one notable example. At the same time, however, there is as yet no agreed explanation for why there is so long, and so variable, an interval between infection with HIV and onset of AIDS. Indeed, I guess that many researchers in this field do not even think about this question. But I suspect the answer may necessarily involve understanding how whole populations of different strains of HIV interact with whole populations of different kinds of immune system cells, within infected individuals. And understanding the nonlinear dynamics of such a system will require mathematical models with similarities to and differences from those that have helped us understand population-level problems in ecology and infectious diseases (8). It may even be that the design of effective vaccines against protean agents like HIV or malaria will require such population-level understanding. As yet, this mathematically theoretical aspect of immunology is even less to be found in textbooks than were mathematical models in ecology texts a generation ago. I venture to predict that the corresponding immunology texts will indeed look different in 20, or even 10, years' time.

In short, mathematical models have proved to have many uses and to take many forms in the life sciences. We all, by this time, acknowledge the usefulness of statistical recipes to help design and analyze experiments. More familiar in some areas than others are the benefits of mathematical studies that underpin pattern-seeking and other software that is indispensable in elucidating genomes, and ultimately in understanding how living things assemble themselves. Very generally useful are still-unfolding advances that illuminate the frequently counterintuitive behavior of nonlinear dynamical systems of many kinds.

Mathematics, however, does not have the long-standing relation to the life sciences that it does to the physical sciences and engineering. It is therefore not surprising to find occasional abuses. Some have been sketched above. Particularly tricky are instances in which conventional statistical packages (often based on assumptions of an underlying Gaussian distribution—the central limit theorem) are applied to situations involving highly nonlinear dynamical processes (which can often lead to situations in which "rare events" are significantly more common than Gaussian distributions suggest) (9). Perhaps most common among abuses, and not always easy to recognize, are situations where mathematical models are constructed with an excruciating abundance of detail in some aspects, whilst other important facets of the problem are misty or a vital parameter is uncertain to within, at best, an order of magnitude. It makes no sense to convey a beguiling sense of "reality" with irrelevant detail, when other equally important factors can only be guessed at. Above all, remember Einstein's dictum: "models should be as simple as possible, but not more so."

References and Notes
1. R. A. Fisher, *The Genetical Theory of Natural Selection* (Dover [reprint], New York, 1958).
2. G. H. Hardy, *Science* **28**, 49 (1908).
3. For a more full discussion see R. M. Anderson and R. M. May [*Infectious Diseases of Humans* (Oxford Univ. Press, 1991), chap. 11].
4. J. Bongaarts, *Stat. Med.* **8**, 103 (1989).
5. R. MacArthur, *Geographical Ecology* (Harper & Row, New York, 1972).
6. For further discussion, see G. Farmelo (Ed.) [*It Must be Beautiful: Great Equations of Modern Science* (Granta, London, 2002), pp. 212–229] and (7).
7. J. Gleick, *Chaos: Making a New Science* (Viking, New York, 1987).
8. M. A. Nowak, R. M. May, *Virus Dynamics: Mathematical Principles of Immunology and Virology* (Oxford Univ. Press, Oxford, 2000).
9. N. Goldenfield, L. P. Kadanoff, *Science* **284**, 87 (1999).
10. M. J. Keeling *et al.*, *Nature* **421**, 136 (2003).
11. I thank many colleagues, particularly S. Levin, A. Grafen, D. Rubenstein, and M. Stumpf for helpful conversations, and the Leverhulme Trust and the Nuffield Foundation for their support.

REVIEW

Evolutionary Dynamics of Biological Games

Martin A. Nowak[1]* and Karl Sigmund[2,3]

Darwinian dynamics based on mutation and selection form the core of mathematical models for adaptation and coevolution of biological populations. The evolutionary outcome is often not a fitness-maximizing equilibrium but can include oscillations and chaos. For studying frequency-dependent selection, game-theoretic arguments are more appropriate than optimization algorithms. Replicator and adaptive dynamics describe short- and long-term evolution in phenotype space and have found applications ranging from animal behavior and ecology to speciation, macroevolution, and human language. Evolutionary game theory is an essential component of a mathematical and computational approach to biology.

Evolution through natural selection is often understood to imply improvement and progress. A heritable trait that confers to its bearer a higher fitness will spread within the population. The average fitness of the population would therefore be expected to increase over time. This is often pictured as a steady ascent on a so-called fitness landscape. The landscape metaphor suggests some solid ground over which the population moves. This paradigm (1), which is also widespread in the theory of genetic algorithms (2), neglects one-half of the evolutionary mechanism: Although the environment selects the adaptations, these adaptations can shape the environment. By moving across a fitness landscape, populations change that landscape (Fig. 1).

This is particularly clear if several populations interact, because each population can be part of the fitness landscape of the other. A host's successful immune response to a pathogen, for instance, will exert selection pressure leading to adapted strains of pathogens, and vice versa (3–5). But even within a single population, the fitness of a trait often depends on the prevalence of that trait: The selective advantage of a given tree height, for example, depends on the heights of neighboring trees. Similarly, the success of a given sex ratio depends on the overall sex ratio in the population.

Therefore, the fitness landscape is shaped by the phenotypic distributions of the involved populations. As the population moves through the fitness landscape, new peaks and valleys form, channeling its further motion. This viewpoint affects not only the intuition of evolutionary biologists but also their theoretical tools. The proper technique for describing uphill motion on solid ground is optimization theory, a set of mathematical techniques developed in the past 300 years, mostly to solve physical or technical problems. If the adaptive steps, however, imply changes in the environ-

[1]Program for Evolutionary Dynamics, Department of Mathematics, Department of Organismic and Evolutionary Biology, Harvard University, 1 Brattle Square, Cambridge, MA 02138, USA. [2]Faculty for Mathematics, University of Vienna, Nordberggasse 15, A-1090 Vienna, Austria. [3]International Institute for Applied Systems Analysis (IIASA), A-2361 Laxenburg, Austria.

*To whom correspondence should be addressed. E-mail: martin_nowak@harvard.edu

ment, eventually necessitating new adaptations, then game theory is the appropriate framework. This technique originated more than 50 years ago to tackle economic and social problems involving interdependencies among several agents. Evolutionary biologists soon understood its potential and started applying it to evolutionary problems (6–8). The success of a strategy in a game depends on the co-player's strategy, much as the fitness of a phenotype depends on the composition of the population. Roughly speaking, game theory is the mathematical toolbox for methodological individualism, the systematic attempt to found social theory on the actions and needs of individual agents (9, 10). For outcomes shaped by "selfish genes" or by the selfish "homo economicus," this is the proper instrument.

Biological Games

The number of papers applying game theory to biological problems is in the thousands (11). The first use of game-theoretic arguments can be found in the field of sex-ratio theory (12, 13). Evolutionarily stable strategies (ESSs) were introduced to discuss the prevalence of ritual fighting in interspecific animal conflicts, which also led to asymmetric and multistep games (7, 14). Evolutions of

root structure or tree height are problems of resource allocation (15). Conflicts concerning mate choice (16), sibling rivalry (17), and parent-offspring antagonism (18) are a rich mine of game-theoretic models; so are social foraging, dispersal, and habitat selection (19). The arms races between predators and prey, or between parasites and their hosts, offer many examples of games between distinct populations (20). Communication in its widest sense, including alarm calls, threat displays, or sexual advertisement, lead to game-theoretic problems concerning bluff and honest signaling (21). Acquisition and performance of human language in a heterogeneous population can be studied as an evolutionary game (22). Increasingly, evolutionary game theory is used in economic and social sciences and applied to experimental games with human subjects (23–27). Even genes, bacteria, organelles, and viruses can be engaged in games of cooperation and conflict (28–31).

The evolution of virulence of infectious agents is another vast field that makes use of game-theoretic arguments. The classical understanding, based on constant selection, is that parasites evolve to maximize their basic reproductive ratio (32). Frequency-dependent selection arises when several parasite strains super-
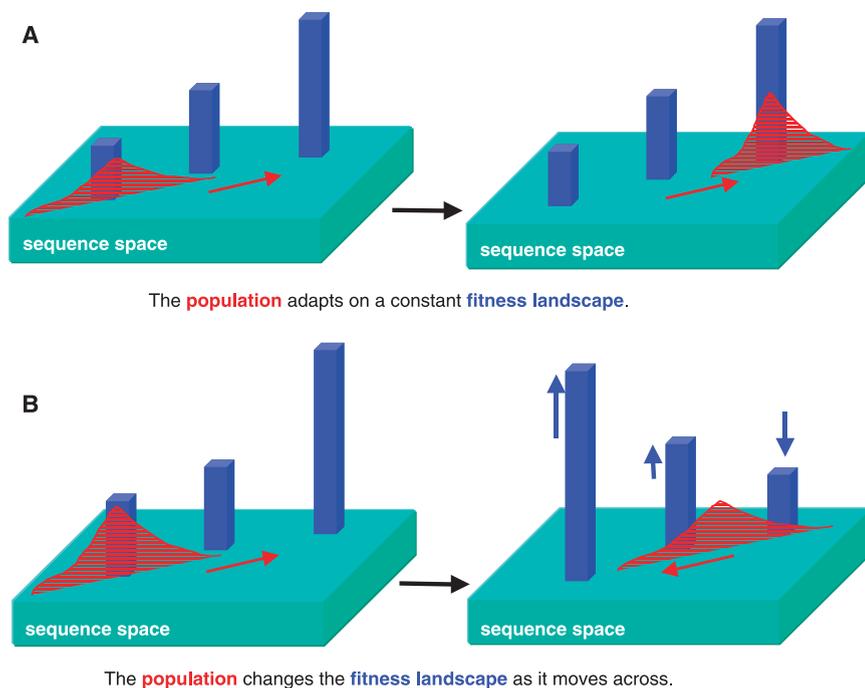
infect the same host or when rapid evolution generates many different parasite mutants in any one infected individual (33). Lack of cooperation among parasites can lead to shortsighted, maladapted levels of excessive virulence harming both host and parasite.

The growth in the range of applications demanded an extension of classical game theory, away from the prevalent static doctrine dominated by the equilibrium notion of Nash and by the quest for a "unique solution" to rational play. The concepts of "unbeatable strategy" (6) and "evolutionary stability" (7) implicitly assumed some underlying population dynamics describing the potential success of invading mutants and, more generally, the interplay of mutation with frequency-dependent selection. An exact formulation of these population dynamics depends on the structure of the population (for instance, well-mixed or sessile), on the mechanisms for the transmission of the relevant traits (by genetic inheritance or cultural learning), and on the time scales underlying the evolutionary and ecological processes. Remarkably, for each of the plethora of conceivable adjustment dynamics (34), there exist simple games (variants of "rock-scissors-paper") where the long-term outcome is not a Nash equilibrium but endless regular or irregular oscillations: Hence, the static approach is in principle unable to provide a full analysis. Only a dynamical theory can describe the "Red Queen" phenomena (35) that are prevalent in evolution, for example, in host-parasite interactions, in the arms races between predators and prey, or in fluctuating degrees of cooperation (36).

Before reviewing the recent developments in evolutionary game theory, we sketch a few basic types of interactions that help to familiarize readers with terminology. A "game" is an interaction between a set of individuals. These "players" act according to their behavioral phenotypes, which are called "strategies." The players' payoffs, which translate into fitness, depend, in general, on their own strategy and on that of their co-players. A tree's height, a parent's sex ratio, a parasite's virulence, a female's choosiness, or a male's ornament are instances of strategies. This terminology is by now well established, but occasionally still induces reactions like "animals don't play games."

Strategic Interactions and Population Structures

The conceptually simplest games offer only two strategies and four outcomes. We illustrate this with the example of some cooperative interaction between two players. Each has the possibility to cooperate, C, or to defect, D. If both play C, both obtain a higher payoff than if both play D, but if the coplayer cooperates, it is better to defect and to exploit

The **population** adapts on a constant **fitness landscape**.

The **population** changes the **fitness landscape** as it moves across.

Fig. 1. A traditional perspective of evolution is adaptation on a constant fitness landscape. Genomes are arranged in sequence space in such a way that nearest neighbors differ in one base substitution. Considering all genomes of length L leads to an L-dimensional discrete lattice. In each dimension, there are four choices representing the four bases. Fitness landscape is a high-dimensional mountain range that assigns each genome a fitness value (reproductive rate). (**A**) A population of genomes (a quasi-species) moves, because of mutation and selection, through sequence space, adapting to a constant fitness landscape. (**B**) In evolutionary game dynamics, the population changes the fitness landscape as it moves through sequence space. In certain games, some strategies do well as long as they are rare but lose out when common. This can lead to stable coexistence or unpredictable oscillations.

the other's efforts. Two alternatives can hold if the coplayer defects: (i) it could be best also to defect, in order not to get exploited; this case represents the well-known prisoner's dilemma game (*37*); or (ii) it could be best to cooperate nevertheless, which leads to the snowdrift game (*38*). If two drivers are caught with their cars in a snowdrift and one of them refuses to cooperate in shoveling a way out, the other driver is better off to cooperate unilaterally, rather than spend the night freezing. In the prisoner's dilemma game, D is a best reply no matter whether the co-player uses C or D. In the snowdrift game, each strategy is a best reply to the other.

The snowdrift game is also known as hawk-dove or chicken game. The first name comes from a situation where animals fight for a territory: "Hawks" escalate the fight, risking serious injury, whereas "doves" flee when the opponent escalates. In the chicken game, two cars are heading for a collision. The loser chickens out, while the winner stays on track. Big-time loss occurs when both stay on track.
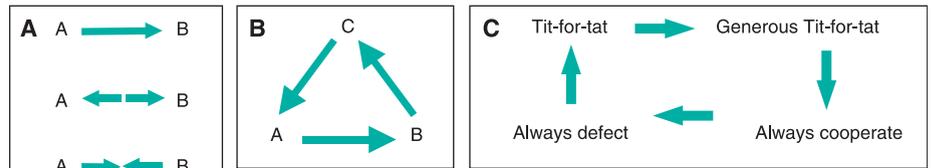
A major challenge in experimental or observational studies is to determine the ranking of payoff values. For example, the interaction of female lions defending a territory against invaders has been interpreted as a prisoner's dilemma (*39*). But, the observation that cooperating and defecting lionesses happily coexist in a group makes it more likely that these lions play chicken.

The payoff values of a prisoner's dilemma have been measured for selection between two mutants of the bacteriophage Φ6. The cooperator builds large amounts of products required for reproduction, whereas the defector specializes in the use of these products when both mutants are in the same cell (*40*, *41*).

A biochemical example for the prisoner's dilemma is provided by the evolution of adenosine triphosphate (ATP)–producing pathways. Cooperators have a low rate but high yield of ATP production, whereas defectors have a high rate but low yield. The resulting game could have played a major role in the emergence of multicellularity (*29*).

In general, evolutionary dynamics of two strategies, A and B, have four outcomes (Fig. 2A). (i) Dominance: A vanishes, if B is the best reply to both A and B. (ii) Bistability: Either A or B vanishes, depending on the initial mixture, if each strategy is the best reply to itself. (iii) Coexistence: A and B coexist in stable proportions, if each strategy is the best reply to the other. (iv) Neutrality: The frequencies of A and B are only subject to random drift, if each strategy fares as well as the other for any composition of the population. The former three cases correspond to the familiar ecological scenarios of two-species competition (*42*).

Examples for all four cases can be found in the repeated prisoner's dilemma, where an interaction between two players consists of many rounds. Tit-for-tat (TFT) is a strategy which cooperates in the first round and then repeats whatever the other player did in the previous round. "Always defect" (AllD) is bistable with TFT if the average number of rounds is sufficiently high. "Always cooperate" (AllC) and TFT are neutral if there is no noise and can coexist in the presence of noise. AllC is dominated by AllD.

With three strategies, the outcome can be considerably more complicated, mainly because of the possibility of "rock-scissors-paper" cycles: It may happen that B dominates A, C dominates B, and A, in turn, dominates C (Fig. 2B). In this case, the dynamics can lead, depending on the relative strengths of the interactions, to either long-term coexistence of all three strategies or to ever-increasing oscillations ending, eventually, with the elimination of two strategies. Examples of such situations, originally viewed as theoretical issues only, have recently been found in nature. (i) There exist three morphs of the male lizard *Uta stansburiana* who differ in their throat color and in their mate-guarding behavior. Type A is monogamous and succeeds in preventing other males from approaching their mate. Type B is polygamous and less efficient, having to split its efforts on several females. Type C does not engage in female-guarding behavior at all but roams around in search of sneaky matings (*43*). (ii) There exist three strains of *Escherichia coli* bacteria. Type A releases toxic colicin and produces, for its own protection, an immunity protein. Type B produces the immunity protein only. Type C produces neither toxin nor immunity (*44*).

We have implicitly assumed so far that interactions within a population are on the basis of random encounters. In many important situations, however, such well-mixing cannot be assumed, and the population structure affects the outcome substantially. If cooperators, for example, preferentially assort with other cooperators, they need not be outcompeted by defectors. Such preferential assortment can be achieved, for instance, if players mostly interact with close relatives (*45*). This raises the issue of kin selection and, more generally, group selection: In groups with many cooperators, average fitness will be higher (*46*). If the degree of positive assortment in the formation of groups is sufficiently high, then cooperators will not be eliminated, although within each group they do worse than defectors. The basic theoretical tool for investigating such "viscous" populations is the Price equation, which describes the growth rate of a strategy as the sum of two terms denoting selection within and between groups.

In a similar way, the outcome of strategic interactions can be affected by preferential assortment with close neighbors. If games are played between neighbors and offspring move only to adjacent sites, then dominated strategies need not be eliminated. The best-known example occurs for the prisoner's dilemma, where unconditional cooperators can subsist, often in frequencies and patterns wildly oscillating in space and time (*47*, *48*) (Fig. 3). A lab experiment has recently highlighted the effect of spatial structure on the population dynamics of the three aforementioned *E. coli* strains engaged in rock-scissors-paper competition. If they live in a well-mixed flask and the resulting population is used to seed, after a few bacterial generations, another flask, then such serial transfers will eventually lead to the survival of a single strain only. If, however, the bacteria grow on the surface of an agar plate and a two-dimensional sample of this surface is used to seed the next agar plate, without altering the neighborhood structure, then the resulting sequence of serial transfers will preserve all three strains (*44*). Often, spatial structure tends to allow more diversity than prevails in well-mixed populations.

Replicator Dynamics and Short-Term Evolution

What are the dynamical systems used to analyze frequency-dependent selection for biological games? A standard tool is the replica-

Fig. 2. (**A**) Evolutionary game dynamics of two strategies admit four cases. B dominates A, A and B are bistable, A and B coexist, and A and B are neutral. (**B**) For three strategies, the possibility arises that A dominates C dominates B dominates A. Depending on the parameters of the game, the evolutionary trajectories can spiral inwards, leading to stable coexistence of all three strategies, or spiral outwards, leading to a random extinction of two strategies. (**C**) In the repeated prisoner's dilemma, there is a fundamental oscillation between cooperation and defection. AllD can be replaced by TFT, which loses to the forgiving and error-prone generous TFT, which is undermined via neutral drift by AllC, which invites invasion by AllD, resembling cycles of war and peace.

tor equation (*49–51*) (Fig. 4). It assumes a well-mixed population with a finite number of strategies and posits that the per capita growth rate of each strategy is proportional to its fitness. If the fitness values are independent of the frequencies, then the average fitness of the population will grow: This yields the usual hill-climbing dynamics and results in the survival of the fittest only. But if the fitness values are frequency-dependent, as is usual with evolutionary games, the average fitness need not grow (*1*). Several strategies can coexist in steady or fluctuating frequencies. Chaotic dynamics are possible.

Many examples of replicator dynamics use linear fitness functions, which is the case if payoff results from pairwise encounters. Games can also be played among larger groups, for instance, public goods games (*52*), or even within the entire population, as in the sex-ratio game. In this case, fitness is usually a nonlinear function of the frequencies.

Replicator dynamics have their basis in the assumption of clonal, genetic reproduction, but these dynamics and variants of it can also describe the spread of strategies by nongenetic mechanisms such as learning, imitation, or other forms of cultural evolution (*50, 53, 54*).

This array of dynamical systems has to be modified if spatial structure is included in the model. If the population is distributed in a one- or two-dimensional continuum, for instance, then reaction-diffusion equations can describe the spread of strategies via traveling waves or the formation of stable spatial patterns (*55*). If players are territorial, some models typically assume that they are located on a grid-like structure and interact with their nearest neighbors only (*47*). Players adopt

the strategy of whichever neighbor obtained the highest total payoff. Again, many variants are possible in terms of lattice geometry, neighborhood structures, and whether updating is synchronous or not. Cellular automata of this type have a rich dynamic behavior, possibly as complex as universal Turing machines. The relationship between spatial models and "mean-field" or "pair-approximations" offer many challenges of mathematical and computational nature (*56, 57*).

Replicator dynamics deal with relative abundances (equal to frequencies). If fitness values depend in addition on the total population size, as is natural in many ecological contexts, then the dynamics have to be supplemented with an equation for the total population size or replaced by equations for absolute abundances (*58*).

## Adaptive Dynamics and Long-Term Evolution

The replicator equation describes selection only, no drift and no mutation. A strategy missing in the initial population remains absent. The dynamics can be modified by adding a small steady rate of miscopying. The resulting outcome will depend on the relative sizes of the mutation rates. A more widespread approach in evolutionary games is to investigate the impact of mutations without explicitly modeling their origin. The new strategy is introduced with an initial frequency that is very small (so that it does not affect the fitness of the residents). Will that frequency decrease or not? Will the mutant be eliminated under the effect of selection, or will it invade the population? To return to the case of two types A and B only: Dominance holds if A can invade B and B cannot invade A (or vice versa), coexistence holds if A can invade

B and B can invade A, and bistability holds if neither type can invade the other.

More generally, given a resident population with its set of strategies, one can study the fate of an additional strategy by considering the replicator dynamics in the augmented space (resident and mutant population) and computing the growth rate of the mutant for small values. If this growth rate, the "Darwinian fitness" of the mutant, is positive, then the mutant will invade. In more complex situations, if the resident population is not at equilibrium, the population is age-structured,

A Quasi-species equation

$$\dot{x}_i = \Sigma_j x_j f_j q_{ji} - \phi x_i$$

B Replicator equation

$$\dot{x}_i = x_i [f_i(\mathsf{x}) - \phi]$$

C Replicator-mutator equation

$$\dot{x}_i = \Sigma_j x_j f_j(\mathsf{x}) q_{ji} - \phi x_i$$

D Adaptive dynamics

$$\dot{p} = \left. \frac{\partial f(q, p)}{\partial q} \right|_{q = p}$$

E Price equation

$$\dot{E}(p) = Cov(f, p) + E(\dot{p})$$

Fig. 4. (A) The quasi-species equation describes deterministic mutation-selection dynamics on a constant fitness landscapes. Genomes are enumerated by $i = 0, .., n$. The frequency of genome i is denoted by x_i, its reproductive rate (fitness) by f_i. The mutation probability from j to i is given by q_{ji}. The average fitness of the population is $\phi = \Sigma_i x_i f_i$. **(B)** The replicator equation describes deterministic but frequency-dependent selection dynamics. The fitness, f_i, of type i is a function of the frequencies of all strategies (phenotypes), $\vec{x} = (x_1, \ldots, x_n)$. For pairwise interactions, it is natural to consider linear fitness functions, $f_i = \Sigma_j x_j a_{ij}$, where the a_{ij} values denote the payoff matrix of the game. **(C)** The replicator-mutator equation combines a mutation matrix with frequency-dependent selection. It can also be used for models of language evolution. Language acquisition is described by the $[q_{ij}]$ matrix. **(D)** Adaptive dynamics describe the evolution of a trait value, p. The function $f(q,p)$ denotes the payoff for a q individual in a homogenous population with trait p. The partial derivative of this function at $q = p$ determines the rate of change of p. **(E)** The Price equation is a general description of evolutionary change of a trait value, p. The population average of the trait value is given by $E(p) = \Sigma_i x_i p_i$ and the covariance between trait and fitness by $Cov(f,p) = \Sigma_i x_i p_i f_i - E(p)E(f)$. The Price equation provides a link between the replicator framework and adaptive dynamics, thereby unifying evolutionary dynamics (*83*).
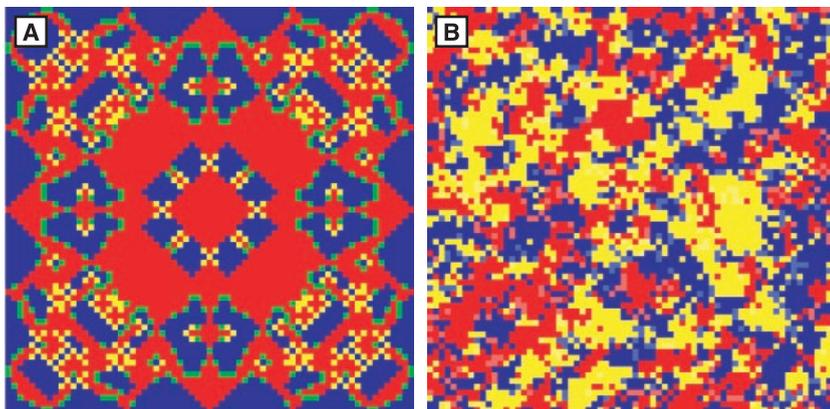
Fig. 3. John von Neuman invented both game theory and cellular automata. Here, the two fields merge. Spatial games of evolution can generate kaleidoscopes, dynamic fractals, gliders, and so on. **(A)** Spatial prisoner's dilemma with $R = 1$, $P = 0$, $T = 1.4$, $S = -0.1$, denoting, respectively, the payoff for two cooperators, two defectors, a defector from a cooperator, and a cooperator from a defector. We use synchronous updating and a symmetric initial condition. Each cell is given to the player with highest payoff in the neighborhood. Color code: red, D; blue, C; green, new C; yellow, new D. **(B)** Rock-scissors-paper game with random initial condition and asynchronous lattice update proportional to fitness (*82*).

or the environment is stochastic, the fitness definition must be adapted (*59*). A successful invasion by itself tells us nothing about the subsequent fate of the mutant. It is possible that (i) it gets added to the resident cast, so that the new population contains one more strategy, (ii) eliminates one or several of the resident strategies, or (iii) eventually gets eliminated itself in a "resident strikes back" reaction (*60*).

The invasibility issue forms the core of evolutionary game theory. A strategy is said to be evolutionarily stable if a resident population cannot be invaded by a mutant strategy under the influence of natural selection, as long as its initial frequency is sufficiently small (*7*). Checking for this condition requires the specification of the range of possible mutants. In the special case that the resident actually dominates all mutants, the strategy is said to be "unbeatable" (*6*). The demands on an ESS are weaker: The mutant must either be dominated or else form a bistable pair with the resident strategy. In this latter case, if the initial frequency of the mutant exceeds a so-called "invasion barrier," it can spread and eventually eliminate the resident. Both stability concepts are variants of the game-theoretic notion of a Nash equilibrium, meaning in this context a strategy with the property that, if all players adopt it, none has an incentive to deviate unilaterally. Whereas Nash equilibria exist for every game, this is not the case for unbeatable or ESSs. It should also be stressed that such strategies are not necessarily optimum: It can be that all players would be better off if they jointly deviated, in a correlated way, with the use of another strategy (for example, cooperation in the prisoner's dilemma).

Such a concerted action is beyond the means of natural selection, but an evolutionarily stable population can conceivably be invaded simultaneously by two or more mutant strategies (*61*).

The theoretical approach through adaptive dynamics considers the opposite, "mutation-limited" scenario: Only one mutant enters at a time, and it is either driven to extinction or fixed in the population before the next mutant occurs (Fig. 5). Thus, mutational steps are supposed to occur only rarely, compared with the time needed to substitute the resident with the mutant, and the population is essentially always monomorphic. The successful invasion attempts lead to a substitution sequence in trait space, that is, the space of all possible strategies, not only those actually present within the population. This phenotypic space is in general a continuum (of possible tree heights, sex ratios, or ages of maturation) and can be multi- and even infinite-dimensional.

If the mutant phenotype is always very close to the resident's, then the discrete-step substitution sequence can be approximated by a smooth orbit of the so-called adaptive dynamics in trait space (*62–64*), pointing into the direction of the most favorable among all possible closeby traits. Whereas the space of replicator dynamics consists of the frequencies of finitely many phenotypes, adaptive dynamics along the local fitness gradient describes evolution in the continuous space of all possible phenotypes. The price for the richer strategy space consists in assuming that the residents are all using the same strategy.

With adaptive dynamics, the motivation tends to shift from behavioral aspects to more ecological issues. Both evolutionary and adaptive dynamics can be used to mod-

el the interaction of several populations. The growth rates of the different strategies will depend, in general, on their frequencies within these populations and on the densities of the populations. The usual, ecological models, for instance, for predator-prey interactions, are limiting cases obtained by assuming that the evolution within each population occurs at a much slower rate than the population regulation. This assumption has recently been challenged in several case studies of rapid evolution showing that population cycles can be driven by evolutionary factors (*65*).

Adaptive dynamics need not lead to an evolutionarily stable state. Cycling can continue forever (*66*). Moreover, if adaptive dynamics converge to an equilibrium, this need not be an ESS. An attractor of adaptive dynamics is said to be a convergent stable strategy (*67, 68*). Conversely, an evolutionarily stable equilibrium can be unattainable (*67, 69, 70*). In that case, if the resident population is at the equilibrium, fitness is maximum and no mutant can invade, but if the resident population is close to the equilibrium, then only mutants can invade that are less close by. Figure 6 shows as an example adaptive dynamics of reactive strategies in the repeated prisoner's dilemma.

So-called pairwise invasibility plots offer a geometric tool for analyzing mutation-limited evolution (*71*). In particular, they show that adaptive dynamics can lead to a fitness minimum, where the monomorphic population becomes unstable and has to split up (*72, 19, 73*).
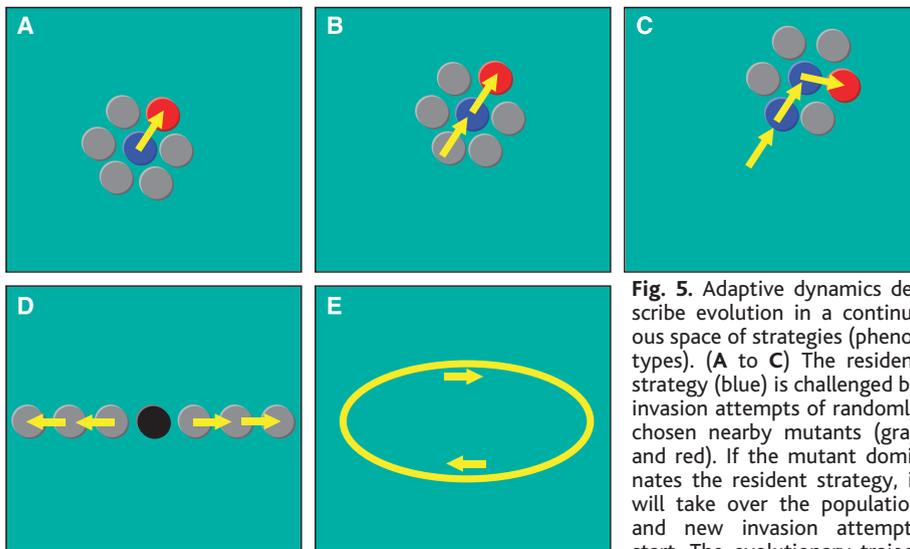
Fig. 5. Adaptive dynamics describe evolution in a continuous space of strategies (phenotypes). (**A** to **C**) The resident strategy (blue) is challenged by invasion attempts of randomly chosen nearby mutants (gray and red). If the mutant dominates the resident strategy, it will take over the population and new invasion attempts start. The evolutionary trajectory (yellow) is most likely pointing in the direction of the mutant (red) that receives maximum payoff from the resident population. (**D**) An ESS (black) can be unattainable for adaptive dynamics. No mutant can invade the ESS, but all nearby mutants are dominated only by strategies that are further away from the ESS. (**E**) Adaptive dynamics can cycle in strategy space.
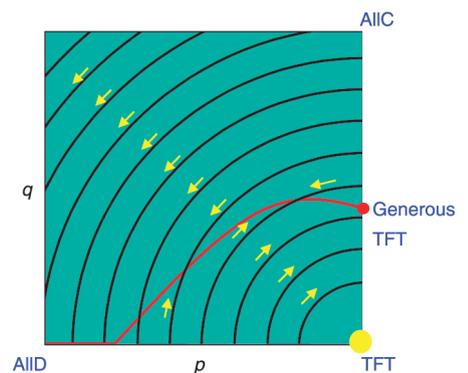
Fig. 6. Adaptive dynamics of reactive strategies in the infinitely repeated prisoner's dilemma. Reactive strategies are characterized by two parameters, p and q, denoting the probability to cooperate after the opponent has cooperated or defected. This set of strategies includes AllD ($p = 0$ and $q = 0$), AllC ($p = 1$ and $q = 1$), and TFT ($p = 1$ and $q = 0$). Adaptive dynamics flow along concentric circles orbiting TFT in the center (*62*). The red curve contains the equilibrium points of adaptive dynamics. They are also Nash equilibria of the game. In the "cooperative region" below the red curve, adaptive dynamics lead to increasing values of p and q. Generous TFT [$p = 1$ and $q = 1-(T - R)/(R - S)$] is the most cooperative Nash equilibrium.

Such evolutionary branching points offer a tempting analogy to sympatric speciation.

A much-studied example involves competition for a common resource, for instance, birds feeding on a type of seed whose size is normally distributed. Without frequency dependence, the birds would be selected to specialize on the most common size of the seed. This specialization, however, might alter the distribution of seed size. Hence, frequency-dependent "disruptive" selection can lead to character divergence, one type specializing on smaller, the other on larger, seeds (*68, 71*).

Evolutionary Games and Population Genetics

Whereas the combination of game-theoretic and ecological models offers no conceptual problems, game theory and population genetics mix less readily. This seems paradoxical: The selection equation in population genetics is a replicator equation, and the alleles are perfect examples of replicators, obeying clonal inheritance. Evolutionary game theory, however, describes evolution in phenotype space. In spite of recent progress in genomics and proteomics, little is known about the genotype-phenotype mapping from a discrete allele space to a continuous trait space.

Under appropriate restrictions (if, for instance, trait space is one-dimensional or if only one locus and at most three strategies are involved), a phenotypic ESS can be viewed as stable end point of a mutation-limited genetic evolution. With recombination between several loci, however, frequency-dependent genetics cannot be fully covered by an invasion analysis at equilibrium points (*74–76*). Genetic constraints as well as linkage and epistasis introduce formidable analytical challenges. Furthermore, the assumption of a monomorphic resident population is obviously a critical issue.

Nevertheless, adaptive dynamics often allow a good understanding of the coevolution of traits under frequency-dependent selection. In particular, branching points of adaptive dynamics may be at the origin of sympatric speciation. In the case of adaptations to different seed sizes, for instance, an interbreeding population can converge to the branching point, and then a mutation for assortative mating based on the trait can invade, eventually splitting the population into two species diverging apart. Speciation through ecological character displacement has long intrigued geneticists (*77*). Numerical simulations show that genes for assortative mating can even spread if they are based on "markers" having nothing to do with the trait value (*78*); on the other hand, assortative mating could fail to exceed an intermediate level if the corresponding genes have small effects only (*79*).

Such examples show that adaptive dynamics can address macroevolutionary issues by analyzing frequency-dependent selection in ecology-driven models. This allows us to study a wide range of nonstatic scenarios, including Red Queen types of cycling and ecological suicide (*80*), which are far removed from the static outcomes suggested by the Panglossian paradigm of hill-climbing in fitness landscapes. Mathematical and computational approaches are increasingly using game theory, rather than optimization, to model the strategic intricacies of coevolution.

Looking Ahead

Evolutionary biology is well grounded in mathematical theory. The way populations change under the influence of mutation and selection can be described by a rich array of mathematical equations. Ideas of evolutionary mechanisms must be formulated in terms of mathematical equations for consistent analysis and meaningful investigation. Most often, Darwinian fitness depends on the relative abundance of individual phenotypes within populations, and therefore game theory is the appropriate mathematical tool.

Many challenges lie ahead. Evolutionary game theory is formulated in terms of phenotypes, thereby ignoring the complexity of the genotype-phenotype mapping. More work is needed on the interaction of strategies encoded in genomic sequences. Most evolutionary game dynamics have been studied in the context of infinitely large populations. We expect that finite population size effects will lead to surprising outcomes and might question the importance of traditional evolutionary stability. Cultural interpretations of replicator dynamics often assume that successful strategies spread by imitation or learning, but the learning of complicated strategies from behavioral observations is a nontrivial task that needs specific investigation. Similarly, studying human language requires a connection between the mathematics of game theory, learning theory, and computational linguistics.

Emerging fields as diverse as metabolic control networks within cells and evolutionary psychology, for example, should benefit from game theory. Metabolic networks respond to incoming signals from other cells and in turn generate outgoing messages, thereby representing a grammar of cellular communication. The performance (fitness) of a particular network depends on networks in other cells, and hence frequency-dependent selection is operating. The main themes of evolutionary psychology include cooperation and communication among individuals and are therefore intrinsically game theoretic. Having a theory of mind, for example, is itself a useful strategy that must help winning certain games.

The applications of evolutionary game theory pervade by now all areas of biology. Interactions among genes, viruses, cells, and humans are often instances of evolutionary games that are amenable to empirical and theoretical investigation. Game theory is the appropriate tool whenever the success of an individual depends on others.

References and Notes

1. S. Wright, *Evolution and the Genetics of Populations II* (Chicago Univ. Press, Chicago, IL, 1969).
2. H. Frauenfelder, A. R. Bishop, A. Garcia, A. Perelson, Eds. *Physica D* **107** (1997).
3. M. A. Nowak, R. M. May, *Virus Dynamics* (Oxford Univ. Press, Oxford, 2000).
4. B. R. Levin, M. Lipsitch, S. Bonhoeffer, *Science* **283**, 806 (1999).
5. S. A. Levin, B. Grenfell, A. Hastings, A. S. Perelson, *Science* **275**, 334 (1997).
6. W. D. Hamilton, *Science* **156**, 477 (1967).
7. J. Maynard Smith, *Evolution and the Theory of Games* (Cambridge Univ. Press, Cambridge, 1982)
8. J. Maynard Smith, G. R Price, *Nature* **246**, 15 (1973).
9. K. Binmore, *Game Theory and the Social Contract* (MIT Press, Cambridge, MA, 1994).
10. D. Fudenberg, J. Tirole, *Game Theory*, (MIT Press, Cambridge, MA, 1991).
11. L. A. Dugatkin, H. K. Reeve, Eds. *Game Theory and Animal Behaviour* (Oxford Univ. Press, Oxford, 1998).
12. I. C. W. Hardy, Ed., *Mother Knows Best: Sons or Daughters? A Review of Sex Ratios. Concepts and Methods* (Cambridge Univ. Press, Cambridge, 2002).
13. R. L. Trivers, *Q. Rev. Biol.* **58**, 62 (1983).
14. A. I. Houston, J. M. McNamara, *Models of Adaptive Behaviour: An Approach Based on State* (Cambridge Univ. Press, Cambridge, 1999).
15. D. S. Falster, M. Westoby, *Trends Ecol. Evol.* **18**, 337 (2003).
16. Y. Iwasa, A. Pomiankowski, *Nature* **377**, 420 (1995).
17. D. W. Mock, G. A. Parker, *The Evolution of Sibling Rivalry* (Oxford Univ. Press, Oxford, 1997).
18. H. C. J. Godfray, *Nature* **376**, 133 (1995).
19. M. Doebeli, G. D. Ruxton, *Evolution* **51**, 1730 (1997).
20. P. A. Abrams, H. Matsuda, *Evol. Ecol.* **11**, 1 (1997).
21. R. A. Johnstone, in *Encyclopedia of Evolution*, M. Pagel, Ed. (Oxford Univ. Press, Oxford, 2002), pp. 1059–1062.
22. M. A. Nowak, N. L. Komarova, P. Niyogi, *Nature* **417**, 611 (2002).
23. M. A. Nowak, K. Page, K. Sigmund, *Science* **289**, 1773 (2000).
24. C. Wedekind, M. Milinski, *Science* **288**, 850 (2000).
25. H. Gintis, in *Game Theory Evolving* (Princeton Univ. Press, Princeton, NJ, 2000).
26. J. Henrich, R. Boyd, S. Bowles, C. Camerer E. Fehr, H. Gintis, R. McElreath, *Am. Econ. Rev.* **91**, 7378 (2001).
27. E. Fehr, U. Fischbacher, *Nature* **425**, 75 (2003).
28. G. Velicer, Y. Yu, *Nature* **425**, 75 (2003).
29. T. Pfeiffer, S. Schuster, S. Bonhoeffer, *Science* **292**, 504 (2001).
30. D. Haig, *Genomic Imprinting and Kinship* (Rutgers Univ. Press, Piscataway, NJ, 2002).
31. D. C. Krakauer, A. Mira, *Nature* **403**, 501 (2000).
32. R. M. Anderson, R. M. May, *Infectious Diseases of Humans* (Oxford Univ. Press, Oxford, 1991).
33. M. A. Nowak, R. M. May, *Proc. R. Soc. Lond. B. Biol. Sci.* **255**, 81 (1994).
34. J. Hofbauer, K. Sigmund, *Bull. Am. Math. Soc.* **40**, 479 (2003).
35. In Alice's Wonderland, the Red Queen has to run to remain on the same spot. For example, hosts have to evolve continuously to maintain the status quo with parasites (*81*).
36. M. A. Nowak, K. Sigmund, *J. Theor. Biol.* **137**, 21 (1989).
37. R. Axelrod, W. D. Hamilton, *Science* **211**, 1390 (1981).
38. R. Sugden, *The Economics of Rights, Co-operation and Welfare* (Blackwell, Oxford, 1986).
39. R. Heinsohn, C. Packer, *Science* **269**, 1260 (1995).
40. P. E. Turner, L. Chao, *Nature* **398**, 441 (1999).

41. S. Nee, *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **355**, 1607 (2000).
42. R. M. May, *Stability and Complexity in Model Ecosystems* (Princeton Univ. Press, Princeton, NJ, 1973).
43. B. Sinervo, C. M. Lively, *Nature* **380**, 240 (1996).
44. B. Kerr, M. A. Riley, M. W. Feldman, B. J. M. Bohannan, *Nature* **418**, 171 (2002).
45. S. A. Frank, in *Foundations of Social Evolution*, J. R. Krebs, T. H. Clutton-Brock, Eds. (Princeton Univ. Press, Princeton, NJ, 1998).
46. E. Sober, D. S. Wilson, *Unto Others: The Evolution and Psychology of Unselfish Behavior* (Harvard Univ. Press, Cambridge, MA, 1998).
47. M. A. Nowak, R. M. May, *Nature* **359**, 826 (1992).
48. A. Sasaki, W. D. Hamilton, F. Ubeda, *Proc. R. Soc. London Ser. B* **269**, 761 (2002).
49. P. D. Taylor, L. Jonker, *Math. Biosci.* **40**, 145 (1978).
50. J. Hofbauer, K. Sigmund, *Evolutionary Games and Population Dynamics* (Cambridge Univ. Press, Cambridge, 1998).
51. J. W. Weibull, *Evolutionary Game Theory* (MIT Press, Cambridge, MA, 1996).
52. C. Hauert, S. De Monte, J. Hofbauer, K. Sigmund, *Science* **296**, 1129 (2002).
53. T. Vincent, B. Brown, *Evolutionary Game Theory, Natural Selection, and Darwinian Dynamics* (Cambridge Univ. Press, Cambridge, 2003).
54. D. Fudenberg, K. Levine, *The Theory of Learning in Games* (MIT Press, Cambridge, MA, 1998)
55. R. Ferriere, R. Michod, *Am. Nat.* **147**, 692 (1996).
56. U. Dieckmann, R. Law, J. A. J. Metz, Eds., *The Geometry of Ecological Interactions: Simplifying Spatial Complexity* (Cambridge Univ. Press, Cambridge, 2000).
57. R. Durrett, *SIAM (Soc. Ind. Appl. Math.) Review* **41**, 677 (1999).
58. R. Cressman, *Evolutionary Dynamics and Extensive Form Games* (MIT Press, Cambridge, MA, 2003).
59. J. A. J. Metz, R. M. Nisbet, S. A. H. Geritz, *Trends Ecol. Evol.* **7**, 198 (1992).
60. S. D. Mylius, O. J. Diekmann, *J. Theor. Biol.* **211**, 297 (2001).
61. R. Boyd, J. Lorberbaum, *Nature* **327**, 58 (1987).
62. M. A. Nowak, K. Sigmund, *Acta Appl. Math.* **20**, 247 (1990).
63. J. Hofbauer, K. Sigmund, *Appl. Math. Lett.* **3**, 75 (1990).
64. U. Dieckmann, R. Law, *J. Math. Biol.* **34**, 579 (1996).
65. A. P. Hendry, J. K. Wenburg, P. Bentzen, E. C. Volk, T. P. Quinn, *Science* **290**, 516 (2000).
66. U. Dieckmann, P. Marrow, R. Law, *J. Theor. Biol.* **176**, 91 (1995).
67. I. Eshel, *J. Theor. Biol.* **103**, 99 (1983).
68. F. B. Christiansen, *Am. Nat.* **138**, 37 (1991).
69. P. D. Taylor, *Theor. Pop. Biol.* **36**, 125 (1989).
70. M. A. Nowak, *J. Theor. Biol.* **142**, 237 (1990).
71. J. A. J. Metz, S. A. Geritz, G. Meszena, F. J. A. Jacobs, J. S. van Heerwarden, in *Stochastic and Spatial Structures of Dynamical Systems*, S. J. Van Strien, S. M. Verduyn Lunel, Eds. [Koninklijke Nederlandse Academie van Wetenschappen (KNAW), Amsterdam, 1996], pp. 183–231.
72. S. A. Geritz, J. A. J. Metz, E. Kisdi, G. Meszena, *Phys. Rev. Lett.* **78**, 2024 (1997).
73. K. Parvinen, *Bull. Math. Biol.* **61**, 531 (1999).
74. I. Eshel, *J. Math. Biol.* **34**, 485 (1996).
75. R. Bürger, *Am. Nat.* **160**, 661 (2002).
76. P. Hammerstein, *J. Math. Biol.* **34**, 511 (1996).
77. M. Slatkin, *Genetics* **93**, 755 (1979).
78. U. Dieckmann, M. Doebeli, *Nature* **400**, 354 (1999).
79. C. Matessi, A. Gimelfarb, S. Gavrilets, *Select. Mol. Genes Memes* **2**, 41 (2001).
80. M. Gyllenberg, K. Parvinen, U. Dieckmann, *J. Math. Biol.* **45**, 79–105 (2002).
81. L. Van Valen, *Evol. Theory* 1, 1 (1973).
82. For an interactive Web page, see www.univie.ac.at/virtuallabs.
83. K. M. Page, M. A. Nowak, *J. Theor. Biol.* **219**, 93 (2002).
84. Support from the Austrian Science Fund WK 10008, the Packard Foundation, and J. Epstein is gratefully acknowledged.

REVIEW

Inferring Cellular Networks Using Probabilistic Graphical Models

Nir Friedman

High-throughput genome-wide molecular assays, which probe cellular networks from different perspectives, have become central to molecular biology. Probabilistic graphical models are useful for extracting meaningful biological insights from the resulting data sets. These models provide a concise representation of complex cellular networks by composing simpler submodels. Procedures based on well-understood principles for inferring such models from data facilitate a model-based methodology for analysis and discovery. This methodology and its capabilities are illustrated by several recent applications to gene expression data.

Research in molecular biology is undergoing a revolution. The availability of complete genome sequences facilitates the development of high-throughput assays that can probe cells at a genome-wide scale. Such assays measure molecular networks and their components at multiple levels. These include mRNA transcript quantities, protein-protein and protein-DNA interactions, chromatin structure, and protein quantities, localization, and modifications. These rich data illuminate the working of cellular processes from different perspectives and offer much promise for novel insights about these processes (*1*).

The challenge for computational biology is to provide methodologies for transforming high-throughput heterogeneous data sets into biological insights about the underlying mechanisms. Although high-throughput assays provide a global picture, the details are often noisy, hence conclusions should be supported by several types of observations. Integration of data from assays that examine cellular systems from different viewpoints (for instance, gene expression and protein-protein interactions) can lead to a more coherent reconstruction and reduce the effects of noise. To perform such an integration, however, we must understand the biological principles that couple the different measurements. In addition, the conclusions of the analysis should go beyond a mere description of the data and should provide new knowledge about the relevant biological entities and processes, ideally in the form of concrete, testable hypotheses.

To answer this challenge, we need to build models of the biological system. A model is a simplifying abstraction. It generates predictions of system behavior under different conditions (as reflected by observations) and illuminates the roles of various system components in these behaviors. We focus on probabilistic models, which use stochasticity to account for measurement noise, variability in the biological system, and aspects of the system that are not captured by the model.

In a model-based approach to data analysis, we start by defining the space of possible models that we are willing to consider. This modeling decision depends on the phenomena we wish to describe and how they are reflected by the observations. We then use a learning procedure to select the model that best fits the actual observations. (Such procedures are referred to by different names in different disciplines, including inference, estimation, reverse engineering, and system identification.) Finally, we use the learned model to reason about the data, make predictions, and glean insights and hypotheses.

An important aspect of model-based approaches is the shift from a procedural methodology to a declarative one. In a procedural method, we focus on the sequence of steps from the data to the conclusions. For example, when relating transcription factor binding sites in the promoter regions

School of Computer Science and Engineering, Hebrew University, 91904 Jerusalem, Israel. E-mail: nir@cs.huji.ac.il