

CHAPTER ONE

GENERAL INTRODUCTION

The ultimate hope of psychology, of biology, of cognitive science is to fully understand the beast, something forbiddingly complex. This study hopes to narrow the search by exploiting what is known of its circumstance.

Sex is a Lamppost

Sex is a cardinal aspect of any species that practices it. There can scarcely be anything more central to a course of evolution than what leads one to reproduce, apart perhaps from what leads it to survive. Evolution ultimately reduces to these two motives, of surviving and reproducing, and neither without the other.

But if we were to cut from the phylogenetic tree the bough nearest to us, say, at mammals, we would see that, whereas there are many ways of surviving, reproduction is done about the same way throughout. This is not to deny the

diversity in mammalian mating systems, but rather to acknowledge that there are palpable similarities.

Nor is this to deny that surviving has its predictable points. Mammals are eaters, and they all eat about the same way, with a toothed mouth that begins whatever chyme it is along a very unsurprising digestive course. But some are big, and some are small; some eat things that run away; some eat things that stay put. Some eat lots at once and rest; others have to eat all the time; others hoard a stash. This is not even to mention the other part of eating: the being eaten. Many animals have much to run away from. Nor is survival all about eating and being eaten. We harvest by day or by night. We walk, swim, fly, or burrow. We live in seas, fields, forests, ponds, treetops, tree trunks, or maybe amidst the roots. We live in packs, clans, troops, flocks, herds, or by ourselves. We come in all sorts.

So what is common to a wolverine and a wildebeest? They have similar digestion, and they have similar sex. Something worth noting about sex is that, whatever the niche an animal is adapted to, the object of sex is unchanged. Very much of the sex lives of bats, on the wing, in their caves, do they share with hippopotamuses. Genetically, they are peas in a pod. And even the situation is very similar in a number of respects. It is anyway that a male and female must unite, and this does not change, even were the hippopotamus to grow wings and twitter about dark caves. Further, this unravels little even beyond mammals to the lot of sexual organisms.

Sex may be a complex, dynamic system, but unlike an ecosystem, with all its whorls and eddies, between changing predators, prey, parasites, and environments, the players of sex remain the same two, and they retain the same

relation between them, with the same dependency. Two sexes makes it much more tractable than other social circumstances. Any system divided this way becomes easier to characterise. If chess had only knights and bishops, and these were constrained by predictable interdependencies, it would be a simpler game. Of especial interest here, this game could be characterised.

Sex has profound influences upon a species' evolution. Sex is predictable. Together these lead to the expectation that something can be said of an organism about which is known only that it is sexual. The aim of this study is to make predictions for the behaviour of organisms that evolve under these predictable circumstances. The strategy taken is to begin with the rules of the matter, and simulate the way an organism might evolve under such constraints. What further patterns can be predicted given the known constraints, in an arbitrary sexual situation?

As a simple example of how knowing the rules of an aspect of sex leads to a simplification of the problem, consider the mere quality of having two sexes. Sex involves combining pairs. Having two sexes constrains the number of allowable combinations to a quarter of what it would be if we were all of one sort, like trees and snails. If there were 1000 snails, each could mate with any other, for 1000,000 combinations (1000×1000). If there were 1000 mice, in contrast, males needing to mate only with females, and vice versa, would constrain the allowable combinations to 250,000 (500×500), a quarter of that of the snails. Further, any skew in the sex ratios would bring this number down. For instance, if 70% of the population is female, the allowable combinations drops by 40,000 ($300 \times 700 = 210,000$), and so forth.

Sex has many aspects, biological and psychological. Some of these can be isolated and varied separately. For instance, the way genetic knowledge is structured is different in most sexual organisms than the way it is in most asexual organisms. This is the topic of Chapter Two on diploidy and allele dominance. In Chapter Three, the matter is the need to couple. Again, this is a feature particular to sexual organisms. In neither case is it obvious beforehand what deeper influence these aspects will have on a species, but both turn out to have profound effects. These are very different aspects of sex, and so Chapters Two and Three are very different. What bring them together are the motive, the method, and the broad notion that sex is a good place to look for constraints on how a species evolves, and so for our ways.

Sex is Social

Sex is essentially multi-agent. It is not something that makes sense of one, or even of two agents, but is an aspect of the community, of interactions between the individuals. For this reason, it is most profitably studied so, within the context of that interaction, and in a manner that embraces the dynamic and complex nature of the matter. Even very simple behaviours of individual agents combine to yield complexity, and may lead to surprising emergent behaviour in the system. The current study uses predictable aspects of sex to plumb for such emergence.

About the Method

The method throughout is to simulate some tangible aspect of sex in a population of agents that are little other than sexual. This means that the model agent cannot be a natural animal itself, for natural organisms are a great deal more than merely sexual. Though, as argued above, sex works in about the same way across species, there are a million other things, at the same time, working differently. Imagine attempting to infer the ways of cottontails from observing those of jackrabbits! Or of humans from quail (e.g. Köksal, et al., 2004). Natural science is forever bound to instances of the phenomena it studies. In comparative psychology, as in comparative biology, this is a difficulty, for there is nothing arbitrary about species. Evolution is the foundation, and this imposes a sort of relativity upon species. Everything is *ad hoc* – a functional adaptation for a given environment and circumstance. The kinds of chemistry do not each have their own periodic table, the way the kinds of biology each have their own evolutionary histories.

Cross-species inferences are drawn, and well that they are, for simulation is prone to the opposite danger, of being incomplete. As an illustrative example, consider Spiderman. The rationale for Spiderman being as strong as he is, is that he has the relative strength of a spider. The author was apparently impressed to learn that a spider can carry several times its body weight, and imagined how strong a man-sized spider would be. Though this seemed compelling as a child, certainly an 80 kg spider could barely, if at all, carry its own body, let alone another. Only if the world had shrunk, rather than the spider distending, might it

have been able to sling our criminally-minded around. The story is intact, though, for in a Spiderman novel, gravity behaves as Stan Lee would have it. It is not a good model, however, for understanding the natural world. The point of this is that any mistake or omission—some very easy to make, as with the relative brawn of big spiders—in a model can throw it. Its accuracy relies on the foresight of the modeller. A natural animal in the real world, in contrast, is sure to be a complete model of what it is, whether the experimenter is aware of every relevant detail or not.

So the natural organism is bound to be too much as a model, the artificial, too little. The study of either, then, obliges the student to qualify every claim with the sort of vessel it was found in. Such and such is true “...of polar bears”, for instance. The advantage the modeller has on this point is to make claims not about species, filled as they are with black boxes, but about what actually went into the model. That is, he can describe a fictional organism in full detail that was designed for the very purpose of studying the phenomenon of interest, and credit those known details for the observed behaviour, rather than a whole, mysterious animal, in whom the detail is incidental. By fashioning the fictional organism to be nothing other than what it is needed to be, it becomes something from which a generalisation can be made.

Further, in the artificial organism, the object of interest can be systematically varied, and accurately, non-intrusively measured. These are luxuries denied to the natural scientist.

Brief Outline

Chapters Two and Three make up the core of this work. Each models a distinct aspect of sex. As such, these two chapters follow distinct courses before regrouping in Chapter Four for a discussion of the conclusions drawn. Chapter One introduces what is common throughout, including the broadest issues, the methods involved, the simulator and the basic agent.

Chapters Two and Three proceed similarly. First, what can be said of the phenomenon as it appears in nature is given. Then its implementation in the model is described, followed by select results from the simulations. Whatever the results, the interesting thing is understanding how it happened so, and what makes it otherwise, so an attempt to analyse the behaviour of the simulations follows. Along the way, other interesting, unexpected behaviours of the system emerged. Attempts to grasp these are given next. Finally, going through these simulations led to several intuitions about ways sexual systems make sense, and about how the aspects investigated here could lead to improvements in artificial evolution. These two chapters each include a section reserved for recommendations for one about to engineer a sexual system.

The Simulation Environment

Simulations are conducted in VUScape (Buzing, 2003), an environment in the JAWAS framework (JAWAS). JAWAS (Java Artificial Worlds and Agent Societies) is an object-oriented, multi-agent systems development framework for implementing artificial societies. The system is implemented in the Java programming language and agents in that system are Java objects. VUScape is based heavily on the SugarScape environment of Epstein & Axtell (1996).

VUScape and JAWAS were chosen after a careful review of several agent-oriented simulation toolkits, and environments within them. Among those considered were Swarm, Evo, A-Scape, RePast, Starlogo, AgentSheets, and MAML (Multi-Agent Modelling Language). JAWAS was ultimately chosen because it was flexible and accessible, and it was based in Java (a pragmatic concern: a more useful skill to acquire). Further, the details of VUScape—concerns like resource management, notions of space, &c.—were nearer to the way I had wanted the environment to be. Though the environment is mostly unchanged for the present work, the agents studied within it were rewritten extensively for the purpose.

VUScape is a virtual world, a 2-dimensional torus-shaped grid, populated by virtual agents all rambling about in search of a limited but replenishing resource, essential for their survival. They can sense a unit of the resource at some distance, and consume it when it is near. They spend their acquired stores of nutrients with each time step in order to maintain their existence, as if fuel or a tax for life. In the simulation, time is discrete. Each cycle allows each agent to

execute, after which the world's resources are replenished. Each simulation begins with a fixed number of individual agents, scattered randomly over a grid of discrete locations, or cells. Each cell can accommodate multiple agents simultaneously. At the beginning of the simulation, each cell is randomly assigned an amount of resource. Consumed resources are replenished at a fixed rate up to a given maximum.

Evolution as a Criterion

Those agents that fail to maintain a store of resources expire. Given that the agents also reproduce, the population is continually being added to, and so the resource will be scarce, meaning that the lot of the agents is one of perpetual competition. Competition assures that the agents are being evaluated upon an implicit criterion for differential survival. At the beginning of a simulation, each initial agent is randomly assigned a value for each heritable trait, defining the analogue of a gene pool. Since offspring inherit select traits from parents, differential reproduction is autocatalytic — the more an agent reproduces, the more potential its code will have for further reproduction. Evolution is expected to occur in such a situation upon each trait that is both heritable and allowed to vary.

Several heritable traits are progressively added to the agents, and the degree that each evolves is measured. The degree to which the trait evolves—the average strength or frequency of the trait after the passage of many generations,

as opposed to the initial average—indicates the propensity of the trait to proliferate under such circumstances. Will the trait evolve when allowed, given the particular circumstances in each simulation, and to what degree? Under what circumstances will the trait evolve otherwise? In addition to mean values, in some parts, especial attention is paid to the dispersion of trait values among agents.

As an example, consider a case in which the agents' metabolic rate is made heritable in the way described. Supposing metabolic rate—the number of units of nutrient metabolised per time step—is initially allowed to vary from 2 to 15, its average value would begin at 8.5. If it evolved to become very low, this would indicate that a lower metabolic rate is adaptive for the species, as would most often be among digital organisms, for whom the need to harvest is as much the modeller's trick as the apple was to Adam.

Some traits differ in frequency, some in strength, and in the most complex model presented in Chapter Three, some traits are varied in frequency and strength independently. Inheritance is treated in full detail in Chapter Two, where it is varied.

The Agents: SugarAgent and Purple

The actors that appear in the various simulations conducted in the course of this thesis vary in several respects. However, they can be characterised, broadly, as two subtypes of the same basic agent type: SugarAgent, in Chapter

Two, and Purple, in Chapter Three. They both exist in the VUScape environment, and feed upon the same resource, abbreviated “Sugar”, following Epstein & Axtell (1996). The ways in which Purple agent is unlike SugarAgent are the subject of Chapter Three, and so are detailed there.

SugarAgent

The four sorts of agent in Chapter Two are all, essentially, the SugarAgent of Buzing (2003), which is a rewriting of the SugarAgent of Epstein & Axtell (1996), with the only difference here being the modifications in the way inheritance is handled.

In each time step, each SugarAgent performs five functions, abbreviated as: Move, Harvest, Metabolise, Reproduce, Die. In the Move function, the agent senses the contents of several cells in each of the four cardinal directions of its current location, and moves to one of them. It moves in accordance with what it senses to be there, given what motivates it. SugarAgent is motivated by Sugar, preferring higher concentrations over lower. This is where the agent is active. Each agent is permitted to behave before anyone moves on to the next step in the execution cycle. In this way, parallelism in agent action is approximated.

In the Harvest function, the agent increments its personal nutrient store, according to the amount found in the cell, modified by the number of agents sharing. Then in Metabolise, this store is decremented by a fixed amount. If the agent was successful in mating during the current round, Reproduce produces its

offspring, and adds it to the agent population. The Die function checks whether the agent still satisfies the requirements for life, and if not, removes it from consideration. Agents can die of starvation or old age.

Purple

The Purple agent is the star of Chapter Three, for which it was created. Its behaviour patterns are controlled by up to 102 simultaneously variable genes, and so it displays wider variability than SugarAgent. Still, these traits are related to their coupling behaviour—what they can sense and what motivate them—beyond which Purple agent is very much like SugarAgent, upon which it was based. In addition to the five functions of SugarAgent, Purple agents perform a sixth at the top of each round, indicating their mood. This is detailed in Chapter Three.

Two Sexes

Purples and SugarAgents each come in two varieties, male and female. In SugarAgent, these are identical, apart from a flag in their code that indicates what sort they are, and throughout Chapter Two, the issue of which is which never arises. To the Purple, however, sex difference is a crucial issue. Though the sexes may begin alike, they are explicitly invited to evolve according to distinct

trajectories, and at that point, especially, it would be misleading to label either as the male or female. As the thing under scrutiny is not maleness and femaleness, but rather a dynamics to be found in a wholly arbitrary 2-sex situation, to call one or the other male or female from the start would be misguided. Still, as they are being allowed to differ, I shall often need to refer to the sexes separately. Call the two sexes of the Purple agent “Reds” and “Blues”. After the experiments are run, it is hoped that Reds and Blues may be interpreted, in instructive ways, in terms of males and females in particular species, but there is anyway no assumption that Reds, for instance, will correspond to the same sex in each species.

Though the emergence of sex-specificity may be influenced, in some species, by factors such as which sex is left to fend for the zygote, or which contributes a larger gamete, assume a species in which all other costs are equal, and with no parental care. In this case, which agent contributes egg and which sperm should not affect outcomes.

We are here interested in the case of a species that has two sexes, and for which fertilization occurs locally. There are sexual species with more than two sexes, or just one (i.e. hermaphroditic species), and there are others for which fertilization occurs at some distance. These impose different demands upon the evolution of a species' behaviour, and so are not captured in the present treatment. However, these are the exceptions. Though such cases are interesting for other reasons, they are less relevant to understanding organisms more like ourselves.

CHAPTER TWO

DIPLOIDY AND DOMINANCE

Natural selection is a conservative sway. By definition, it consumes variability. It is the favouring of some portion of a distribution of characters over others, such that subsequent generations manifest a range that is, whatever else, shorter than it had been. By this successive shortening of possibilities, if left unchecked, natural selection will eventually dispense of all variability, and the evolving entity—whether species, solution, or programme—will have converged. That is, a population of variants will have been reduced to uniformity.

Though convergence is the aim in most applications of evolutionary computing, it does mean the end of evolution, and so a commitment to a particular variant. The obvious risk is that the chosen variant will be other than optimal. Evolutionary algorithms, particularly those that rely exclusively on local operators, as in natural evolution, carry the risk of converging on a solution that, whether a local optimum or not, is inferior to alternatives elsewhere in the search space. In such cases, convergence is said to be premature. Premature convergence is described as “the preeminent problem” in genetic algorithms

(Potts, Giddens & Yadav, 1994, p.73), and this sentiment resounds in the literature, not only for genetic algorithms, but for evolutionary algorithms in general (Mauldin, 1984; Lin, Punch & Goodman, 1994; Herrera & Lozano, 1996; Leung, Duan, Xu, & Wong, 2001).

Evolutionary algorithms reliably reach optimal solutions on the condition that sufficient variability is maintained (Potts, Giddens & Yadav, 1994). This point has been demonstrated mathematically for some specific evolutionary algorithms (e.g. Rudolph, 1994, for the (μ, λ) -ES, with a class of convex fitness functions; others in Fogel, 2000), and more recently, for the general case, by analysing convergence in an abstraction of the evolutionary algorithm (Leung, Duan, Xu, & Wong, 2001).

The basic approach of dealing with premature convergence by introducing explicit measures to maintain diversity was incited by Mauldin (1984). Mauldin's (1984) solution was to enforce a uniqueness condition on new births: only individuals sufficiently unlike all existing individuals may enter the population. Though this is made less attractive in its requirement for a global operator—both expensive computationally, and cumbersome theoretically—it succeeds in its aim to stave off convergence. Later work corroborates this basic strategy (e.g. Shimodaira, 2000).

Nature has its own mechanisms for assuring diversity. Species do not converge, after all, despite the passage of countless generations. An invariant gene pool would not allow a species to evolve, and so leave it vulnerable to potential environmental change (Smith, 1988), or exploitation by parasites (Ridley, 1995). For natural selection to continue to operate, it requires

something to provide it with a stream of variability. Mutation does this, but inefficiently. As mutation is nothing other than random copying errors, it is more disruptive than constructive. Hence, increasing mutation rates can be counter-productive, as it blindly unravels what benefit evolution has won (Mauldin, 1984; Potts, Giddens & Yadav, 1994; Herrera & Lozano, 1996). Therefore, mutation is constrained to be slow for practical concerns, and is so in natural organisms (Smith, 1988).

Apart from creating diversity, as mutation does, a second effective strategy for avoiding genetic stagnation is to merely slow selection down (Lin, Punch & Goodman, 1994; Gao, 1998). The longer the present variability can be held on to, the more thorough will the search have had the potential to be upon convergence. If the processes that deplete variability can be slowed, selection is allowed to proceed without unravelling its spoils, unlike with mutation. Sexual organisms have such mechanisms, by which they retain variability, allowing even a slow stream of it to accumulate.

Here we consider two ways natural sexual organisms preserve variability by maintaining a reservoir of unused genetic material. Natural selection can only kill the variability that it can see. Genes that have no effect on their possessors will not be systematically selected for or against. Unexpressed genes, even those which are harmful when expressed, remain in circulation in the gene pool. Hence, mechanisms that allow genetic material to hide silent in the genome shield variability from the conservative pruning of natural selection, and so promote future evolution.

One such mechanism is diploidy. The haploid genome, as a data structure, is a list of values, one for each trait, that determines the particular qualities of the individual. The diploid genome is a double list, containing two values for each trait, and the possessor's qualities are determined by some function of these two corresponding values. This means that there is a dissociation between the heritable code (genotype) and the way it is expressed (phenotype). This allows some genetic material to remain silent, and so be passed on whatever its effect might have been had it been expressed. Thus, the diploid genome cautiously preserves genetic variation in a way that the haploid genome does not. Regarding the gene pool as a sort of implicit knowledge, the effect of diploidy is to retain a morsel of memory of the search history (Branke, 1999; Goldberg & Smith, 1987; Ryan, 1996).

Another such mechanism is allele dominance (Campbell, Reece & Mitchell, 1999). Each parent provides an allele, a corresponding version of each trait value. In some cases, only one of these alleles is expressed, and the other lies dormant as if in deference: the dominance of one allele over the other is complete. In other cases, the two alleles combine to express an intermediate quality. This is known as codominance. An important difference is that complete dominance shrouds one allele, letting it remain despite the tests of natural selection, in a manner that codominance does not. A further important difference lies in the way these functions bias inheritance: complete dominance passes on the most radical interpretation of the parents' genes; codominance, the most conservative.

Here we report upon a comparison of the courses of evolution of four versions of an arbitrary sexual species, the sugar agent, a denizen of the JAWAS agent oriented simulation toolkit (JAWAS). Agents inherit traits according to one of two models of genetic transmission—the haploid versus the diploid genome—and with either complete allele dominance (“Mendelian dominance”) or no dominance at all (“codominance”).

Existing Evolutionary Computing Literature on Diploidy and Dominance

Work On Diploidy

Many multi-agent systems rely on evolutionary processes. In many of these, the model of reproduction is sexual. Apart from most fungi and some algae, almost all sexual organisms are diploid. However, the haploid pattern of inheritance is seen in the artificial agents of most of the multi-agent simulation packages available for social science research, including JAWAS (JAWAS), A-Scape (A-Scape), and others. Similarly, sex in artificial agents often ignores the issue of allele dominance. These mechanisms do appear in the artificial life literature, for example, when the explicit object is to model some aspect of sex (e.g. Jaffe, 1996). However, given the ubiquity of these mechanisms in natural organisms, and their significance to a course of evolution, they are bound to be relevant aspects of any model in which reproduction is sexual, and it makes sense to wonder what sort of implicit decisions one is making in leaving them

out.

In general, haploidy is the norm in evolutionary computing. Most applications of genetic algorithms (Kurup; dos Santos, 1996), genetic programming (Vekaria & Clack, 1997), and evolutionary strategies rely on a haploid model of genetic transmission, and it is this that is given as standard in textbook introductions (e.g. Langdon & Poli, 2002; Russell & Norvig, 2003). This model has been set upon very many problems with success, both derived and applied (Russell & Norvig, 2003). When the solution sought is not apt to change, and the search space is constrained, as with many situations, haploidy is likely to do well, and the redundancy in diploidy will offer little benefit (Kurup; Yilmaz & Wu, 2002). Hence, attempts to apply diploidy to such problems have not always found it to offer an improvement (Branke, 2001; Schafer, 2003).

The motivation¹ for implementing diploidy in genetic algorithms (GAs) came from an attempt to deal with dynamic, time-varying problem environments, on which typical GAs fare poorly (Goldberg & Smith, 1987; Singh, 2002; Simoes & Costa, 2002). Because diploidy holds on to its variability, it will be more able to adapt to changing demands. If it has begun to settle on a solution, and the goal changes, a haploid GA may have squandered its only potential to adapt, and be stuck in an outdated optimum (Mauldin, 1984). Goldberg and Smith (1987) demonstrated that diploidy conferred a GA with a sizeable benefit, over a regular haploid GA, on such a problem. Since then, supplementing GAs with diploid representations has been done with similar motives, and the basic findings in Goldberg & Smith (1987) have mostly been

¹ Other researchers choose the diploid model uncritically, for the sake of fidelity (e.g. dos Santos, 1996).

replicated (e.g. Kurup), with various applications (e.g. Osmera, 2003); with variations on the implementation of dominance (e.g. Ryan, 1996; Ng & Wong, 1995; Lewis, Hart & Ritchie, 1998); as well as with variations on the diploid representation itself (e.g. Dasgupta & McGregor, 1992). For an exception, see Yilmaz & Wu (2002). The literature on the use of diploidy in genetic algorithms is concisely surveyed in Branke (1999) and in Singh (2002). Kursawe (1991; cited in Ghosh & Dehuri, 2004) is an example of diploidy applied to evolutionary strategies. The overall conclusion to be drawn from this work is that, in problems that change over time, such as when the fitness function is made to oscillate, diploidy follows the moving target, whereas haploidy fails (Goldberg & Smith, 1987; Schafer, 2003; Osmera, 2003). Survival, to a line of genes, is such a problem.

Evolution for genetic algorithms and evolutionary strategies is unlike that in our multi-agent setup in some relevant respects. To begin, there is no explicit fitness evaluation function here, as in GAs. Fitness in an agent world is, as with natural organisms, implicitly determined by many factors in the environment, and can very feasibly change its demands, for instance, at different population densities. So, though not as much as a line of genes, an artificial lineage may find itself having to deal with oscillating fitness demands, the sort of problem in which diploidy most confers an advantage in GAs. Though fitness in the present study is fairly steady, it is somewhat noisier than the fitness function of a typical GA.

Another reason multi-agent methods might get more out of diploidy than do most genetic algorithms lies in the frequency of fitness evaluations. In our

artificial world, and others like it, fitness evaluation is continuous. In contrast, fitness evaluation in genetic algorithms is (usually) as rare as once per generation. The especial relevance of this to diploidy and dominance lies, again, in its carriage of unexpressed alleles. As these are unexpressed, they are exempt of evaluation, and this exemption lasts for the entire generation. The more fitness evaluations a quality is allowed to pass dormant, the better the genome can hold on to its variability, even while exposed qualities are under harsh selection pressures. In this way, frequent fitness evaluation heightens the effect of genes being permitted to hide in diploidy. This point was clarified empirically in Schafer (2003). Schafer (2003) varied the number of fitness evaluations per generation in a diploid GA. When fitness was evaluated only once each generation, diploidy offered no advantage. When fitness was evaluated twice, a stark advantage was seen. It may not, then, be appropriate to generalise to multi-agent systems, conclusions drawn from work on GAs.

Work On Dominance

As with diploidy, a literature utilising the notion of dominance in genetic algorithms seems to have begun with Goldberg & Smith (1987). Already the conception of dominance had gone beyond that in the present study, with alleles that remember their dominance values. Vekaria and Clack (1997) is an attempt to apply a similar approach to dominance with genetic programming. Following Goldberg & Smith (1987), various dominance schemes were studied in (diploid)

GAs. Some studies developed the representation scheme used (e.g. Ng & Wong, 1995; Schnier & Gero, 1997). The next step was to relax dominance, to begin to allow for incomplete dominance (Ryan, 1996). Then, various schemes by which dominant-recessive relations undergo change were developed (Yilmaz & Wu, 2002; Ng & Wong, 1995; Lewis, Hart & Ritchie, 1998). Dominance schemes are reviewed in Singh (2002). As with diploidy, the bulk of this work has been done with GAs, and so attempts to generalise therefrom should be made with caution.

Work on diploidy and dominance are interwoven. Of the studies mentioned in the previous paragraph, all except Vekaria & Clack (1997) use a diploid representation. Indeed, there are reasons to think that they would work together, and they do largely concur in nature. However, this is all the more reason to pull these two mechanisms apart, and vary them independently of each other, as the current study attempts to do.

Mendelian dominance, as modelled here, is the standard among evolutionary algorithms not explicitly studying dominance. Other studies that approach the issue of allele dominance take the notion beyond the simple dominance scheme modelled here. By our model, a gene that dominates in one case might be recessive in the next. By other models, in contrast, particular alleles are consistently dominant or recessive. Though this is biologically accurate, it skips a computational step, which turns out to have interesting effects of its own. What seems not to have been explicitly modelled and tested before is the entire absence of dominance, as in the case of our codominance condition. The partial dominance in Ryan (1996) is not codominance. It does not return intermediate phenotypes, for instance. And like other dominance

schemes, it is far more complex and derived than the implementations in our model.

We have several requirements of our model of diploidy and dominance, beyond what previous work has provided. We require that it be agent-oriented. Sex is a multi-agent process in essence: it is something that only works in societies of interacting agents. Further, it must be true to nature, and as simple as these processes can be formulated. It must be able to dissociate dominance effects from ploidy effects. And it must be subjected to experiments designed expressly for the sake of studying these effects. We report on such a model here.

Haploidy and Diploidy

Almost all of the variability between individuals of sexual species is due to the shuffling of genes in meiosis (Campbell, Reece & Mitchell, 1999). Meiosis is the process by which a diploid cell divides into haploid cells. Diploid and haploid refer to the amount of genetic material a cell contains. A haploid cell contains a single complement of genes; a diploid cell contains two. Most species that exclusively reproduce asexually have only haploid cells at all stages. In animals that reproduce sexually, each parent contributes, via meiosis, a haploid sex cell (gamete), and these fuse to form a diploid cell (zygote), which is what thereupon grows into the new individual. Its genome is thus diploid. This is the most familiar form that sexual cycles take, but other sexual organisms, notably fungi (Campbell, Reece & Mitchell, 1999), do it slightly differently. In such organisms,

sexual reproduction follows an analogous passage through diploid and haploid phases, except that it is the haploid cell that develops into the mature, multicellular individual. We abbreviate these two varieties of sexual cycle as the diploid pattern, and the haploid pattern, respectively. In the diploid pattern, the genome of the mature organism is diploid, and the seed or egg is haploid; in the haploid pattern, the genome of the adult is haploid, and the dormant unicellular seed-like phase is diploid.

The difference between these, though subtle, is important: it is the mature organism that must contend with the world—that is to do well or poorly in it—and so upon which natural selection may operate. By the haploid pattern, the expression of one set of genes determines the course of evolution; by the diploid pattern, the expression of two sets of genes does. This has important repercussions, for it allows a gap to appear between what is passed on to the next generation (the genotype), and what is expressed (the phenotype). This allows an allele to remain effectively hidden in the genome, where natural selection cannot act upon it.

The haploid genome keeps just one set of instructions for each trait, whether that is the mother's copy, the father's, or some function of the two. This directly determines the way the trait is expressed, and so the differentiation between genotype and phenotype is blurred. Selection is operating directly upon the trait that earned the selection advantage. There can be no mistakes. If an individual does well on account of a high score on some trait, it will have nothing other than that same high score to pass on to its offspring. No low scores can get through. It can be said that selection is transparent under these circumstances.

In contrast, by the diploid procedure, there is a clear distinction between the genotype and the corresponding phenotype. The diploid genotype consists of two sets of instructions, one from each parent. The phenotype is again determined by the genotype, but each trait being a function of two values is, of course, less direct than it being determined by a single value, as in haploidy. So when a successful agent reproduces, its offspring may well inherit a quality very unlike that expressed in the parent. It is this gap between genotype and phenotype in diploidy that allows it to smuggle unused genes past natural selection.

The algorithms we implemented for haploid-style and diploid-style inheritance are described below. In both haploidy and diploidy, the parents' genotypes are taken as the input to the reproduction function. A genotype and a phenotype are returned as outputs.

The Haploidy Algorithm

The genome of the haploid agent consists of a list of values, one for each heritable trait. Each value is a real number, between 0 and 1, expressing a probability to behave in a prescribed manner. Recombination, to an agent that is both sexual and haploid, consists of combining such a list from each parent into one, and assigning the resulting array to the offspring as its genotype.

Inheritance for such an agent can be described in two steps:

1. For each trait, take the corresponding gene from each parent's genotype.
2. Some function, f , of these two values is assigned to the offspring. In principle, if not in nature, the sort of relation designated by f is not constrained. The resulting value determines both how the trait is expressed in the agent (i.e. its phenotype), and also what the agent ultimately passes down to its offspring (i.e. its genotype). For trait X :

$$\text{genotype}[X] (= \text{phenotype}[X]) = f(\text{father.genotype}[X], \text{mother.genotype}[X])$$

The Diploidy Algorithm

The diploid genome consists of a 2-dimensional array. It is a list of pairs of genes. Each heritable trait is represented by a real number for each of two genes, organised in a 2-element array. Inheritance for a diploid agent is a matter of combining the mother's array and the father's array into that of the offspring.

This can be described in three steps:

1. For each trait, take one value, randomly selected, from the appropriate array in each parent's genome. This is the analogue of meiosis.

2. The new agent likewise inherits a 2-element array for each trait. This is filled with one gene from each parent. This is the analogue of fertilisation. For trait X:

$$\text{genotype}[X][0] = \text{mother.genotype}[X][0] \text{ OR } \text{mother.genotype}[X][1]$$

$$\text{genotype}[X][1] = \text{father.genotype}[X][0] \text{ OR } \text{father.genotype}[X][1]$$

3. Some function, f , of the two values in the genotype determines how the trait is expressed in the agent (i.e. its phenotype). For trait X:

$$\text{phenotype}[X] = f(\text{genotype}[X][0], \text{genotype}[X][1])$$

Allele Dominance

There could be any number of functions for determining how the parents' genes combine to give the offspring's phenotype. One item that differentiates such functions in carbon-based organisms is the issue of trait dominance. In some cases, a disagreement between two genes trying to manifest different values for the same character is resolved by compromising upon an intermediate value. In other cases, one allele dominates over the other. When this happens, only one allele is expressed; the other, recessive allele is carried silently. It is this silent allele by which dominance shields genetic variability from natural selection.

Traits differ in the degree to which alleles dominate over others, from complete dominance, to codominance. Each of these suggest a function for determining trait expression — the f in the above algorithms. By simulating both of these extremes, we attempt to capture this dimension. These functions differ in the way they bias inheritance: complete dominance passes on the most radical interpretation of the parents' genes; codominance, the most conservative.

Note that our model does not include alleles that are consistently dominant or recessive to each other, as in natural organisms. In our model, which value dominates is decided anew with every birth. In pea plants, in contrast, the allele for purple flowers always dominates over the allele for white flowers. This omission makes a difference in the degree that certain genes are shielded. Since consistently recessive alleles are only exposed to selection pressures under specific circumstances, such as when two come together, this

insulating effect is heightened as the allele becomes rarer. This means that such alleles are very resistant to extinction. Thus, the results of our model should be regarded as conservative.

Mendelian Dominance

Some traits are all-or-none. If pea plant A has purple flowers, and pea plant B has white flowers, their offspring would express one of these petal colours, rather than a mixture of the two. This is the famous discovery of Gregor Mendel from 1865 (cited in Campbell, Reece & Mitchell, 1999). In our simulation, under this sort of trait expression, agents express only one of the values that represent each trait: either that inherited from the mother, or that inherited from the father. Hence:

$$f(x, y) = x \text{ OR } y$$

Codominance

Some traits are the culmination of very many genes. Some of the relevant genes that the mother contributed will dominate; some of the father's will. Overall, the child's phenotype is expected to approach the average. If agent C has red petals, and agent D has white petals, and if they are a certain species of

snapdragon, all their offspring will have pink petals (Campbell, Reece & Mitchell, 1999). This sort of trait expression is simulated here by simply taking the average of the two values that make up the genome for that trait. That is:

$$f(x, y) = (x + y) / 2$$

The Simulation Environment and Agents

Simulations were conducted in VUScape (Buzing, 2003), an environment in the JAWAS framework (JAWAS). The genome of agents in VUScape consists of two traits: Talk and Listen. The simulation is set up to encourage cooperation among the agents. To harvest the resource found in a given cell, there must be enough agents present, working in tandem. Depending on the quantity of the resource present, to harvest a cell might require the combined efforts of up to four agents. To such an agent, “talk” is a plea for help, and “listen” is the act of responding to another agent’s plea. Each agent has a Talk value, which indicates the probability that the agent will issue a plea for help whenever it happens upon a cell that is too big for it to harvest alone. Each agent also has a Listen value, which indicates the probability that it will preferentially move to a cell in which another agent has talked. Under such circumstances, natural selection favours individuals that talk and listen more frequently (Buzing, 2003; Buzing, Eiben & Schut, 2005). This is how communication works in VUScape.

For the experiments reported herein, we supplemented the reproductive

function of VUScape agents with the methods for genetic transmission described above. The implementation of the haploid-codominant condition (Haploid-CoD) was already present in JAWAS. VUScape's default settings were used, with one exception: we raised the resource growback rate (from 1 to 1.3), so that simulations would support larger populations of agents. Each run was initialised with 1000 agents in a 50 x 50 cell grid. Simulations were run for 3000 time-steps each.

Design

Two criteria are measured for comparing populations: 1) the average values of each trait, Talk and Listen; and 2) the variance seen in these values. These two numbers represent the way each trait is distributed at each time-step over the course of each simulation. In this way, the behaviour of these two dependent measures for each trait is a good description of how the different models of genetic transmission implemented effect evolution.

Simulations were run with one of two implementations of genetic transmission (haploid versus diploid), and with one of two functions for trait expression (Mendelian versus codominant). These are our two independent variables, for a 2 x 2 experimental design. This yields four sorts of populations: haploid and Mendelian (Haploid-Or); haploid and codominant (Haploid-CoD); diploid and Mendelian (Diploid-Or); and diploid and codominant (Diploid-CoD).

Each of these population types were run 15 times (4 of the Haploid-CoD runs were excluded from analyses, on account of early extinction). Each simulation was continued for 3000 time-steps. The mean and variance of each trait (Talk and Listen) were sampled at every 100 iterations.

These data are subjected to two sorts of repeated-measures, 2-factor ANOVA. The factor of primary interest is the sort of population, whether Haploid-Or, Diploid-Or, or Diploid-CoD (3 conditions). The haploid-codominant condition is excluded from analyses on account of the much lower variance in this group — approaching, sometimes reaching, zero. Its inclusion would muddy the statistics used, and unnecessarily, for with no variance, one does not need a statistical test to say that it is different! This condition is highly dissimilar from the other three conditions on all measures.

The other factor is time. When looking at trajectories of evolution over time, or interactions of conditions with time, the entire trend, from initialisation to end, is analysed. Such tests include 31 time-steps (time-step 1, then every 100 iterations). Where the interest is to analyse the end results of evolution, after populations have somewhat settled, only data from the last 6 time samples— from time-step 2500 to 3000—are considered. In previous work with JAWAS (e.g. Buzing, Eiben & Schut, 2005; Eiben, Nitschke & Schut, 2005), simulations were run for 2000 time-steps. Our simulations were run for 3000 time-steps. Bonferroni post hoc tests are used, as a standard test for post hoc effects.

Results

Both of the heritable traits in our agents' genomes were unambiguously advantageous to their survival. Average trait values were expected to increase over generations. But would they do so at the expense of variability? Figures 2.1 and 2.3 show the evolution of Talk, one of the two traits allowed to evolve in the simulations. At the beginning of each simulation, each agent is randomly assigned a number between 0 and 1, which indicates its Talk value. Random assignment of trait values assures two initial conditions: that the average trait value will be roughly 0.5 (Figure 2.1); and that variance will be artificially high (Figure 2.3). So the action of evolution over the first several hundred time-steps can be thought of as correcting these impositions. How do these different algorithms make this correction?

First note that the algorithm followed in the Haploid-CoD condition does not correct it. It rather consolidates the population's initial average trait values. If these values are good, all is well, but in the present case, agents clearly could have done better. Within the first few hundred time-steps, all agents in the simulation have identical trait values. Traits converged to as low as 0.456 and only as high as 0.577, over 11 simulations (Listen: $M = .515$, $SD = .034$, ranging from .456 to .577; Talk: $M = .512$, $SD = .031$, from .471 to .556). Evolution, for these populations, acts to stifle evolution! This cannot be an apt model for an evolutionary process. However, let us note that despite its poor performance, it does initially strike one as somewhat accurate. Indeed, this is the only method of genetic transmission that we took unaltered from the JAWAS simulation




Figure 2.1. Mean Talk values for each condition, over time.




Figure 2.2. Mean Listen values for each condition, over time.

framework. It is in retrospect that we see how it necessarily kills variability. Its comparison with the other three methods is instructive, for it accentuates what is happening differently there.

All of the other three methods allow natural selection to act, and directional evolution is seen to occur. Given that both of the traits that may be so affected are set up to be advantageous to the agents, we expect average values to increase over many generations. This may be regarded as an implicit goal of each of our algorithms: to evolve better harvesters, who use Talk and Listen more often. This clearly occurs in all of these three methods, as can be readily seen in Figures 2.1 and 2.2.

Our primary concern, however, is how each method performs on a second implicit goal: retaining variability. This conflicts with the pressure to score highly, for the nearer agents come to some ideal value, the less variable will the population become with respect to that trait. If a population would evolve such that all agents come to have the same top value, it may be said that they are more adapted to their environment, but the species would be vulnerable to potential environmental change (Smith, 1988), or exploitation by parasites (Ridley, 1995). How will each model of genetic transmission resolve these conflicting demands?

Means

A progressive increase of Listen and Talk values was seen in all Haploid-Or, Diploid-Or, and Diploid-CoD simulations (see Figures 2.1 and 2.2). Since the first generation agents are assigned trait values randomly from 0 to 1, as each simulation begins, the average score is roughly 0.5 (in our 56 simulations, for Listen, $M = .498$, $SD = .030$). At the end of the simulations, with the passage of 3000 time-steps, mean Listen values had increased to 0.965 for Haploid-Or populations ($N = 15$, $SD = .025$, from .898 to .993); 0.929 for Diploid-Or ($N = 15$, $SD = .044$, from .847 to .996); and 0.926 for Diploid-CoD ($N = 15$, $SD = .032$, from .879 to .996). The corresponding increase in Talk is charted in Figure 2.1. This rise lagged behind that of Listen, on account of the nature of these capacities, but was no less consistent. At time-step 3000, the average Talk values had risen to 0.829 for Haploid-Or populations ($N = 15$, $SD = .071$, from .716 to .923); 0.799 for Diploid-Or ($N = 15$, $SD = .064$, from .668 to .893); and 0.743 for Diploid-CoD ($N = 15$, $SD = .068$, from .638 to .872).

A repeated-measures, 2-factor ANOVA (31 time-steps \times 3 conditions) for each trait showed main effects and interactions around. Mean values of the two heritable traits increased over generations (Listen: $F(30) = 886.7$; Talk: $F(30) = 170.5$, p 's $< .001$). The interaction between population type and time proved significant as well (p 's $< .001$), meaning that traits developed according to dissimilar trajectories over time. For both traits, the highest values were obtained in the haploid-Mendelian condition. Furthermore, these heights are achieved more quickly than in other conditions. This is starker in the case of

Listen (Figure 2.2), but apparent in Talk as well, as seen in Figure 2.1. That groups ultimately reached dissimilar heights was confirmed with a 2-factor ANOVA (6 time-steps \times 3 conditions) for each trait, on the latter portion of the simulations (Listen: $F(2) = 7.09$, $p = .002$; Talk: $F(2) = 8.16$, $p = .001$). Post hoc tests show that Haploid-Or populations evolved notably higher Listen scores than both diploid conditions (p 's $< .02$), which were similar to each other ($p = 1.000$). A similar pattern is seen in the evolution of Talk, except that the diploid-Mendelian condition rose to levels nearer to the haploid-Mendelian condition ($p > .5$), and both of these were significantly higher than the diploid-codominant (p 's $< .05$). Populations that inherited Talk according to the Mendelian function evolved to possess higher trait values than those using the codominance model, though this was not seen in Listen. All inferential statistics are summarised in Table 2.1.




Figure 2.3. Mean variance of Talk values for each condition, over time.




Figure 2.4. Mean variance of Listen values for each condition, over time.

Variance

In each condition, as evolution proceeds and trait values go up, variance goes down. Agents become more homogeneous. This is expected, but to what extent will each algorithm relinquish variability for high scores? Two differences can be seen between the trends plotted in Figure 2.3 and in Figure 2.4: the rate of their decline; and the extent that they dip. In both codominance conditions variance drops very rapidly. This is striking in comparison to the much more gradual declines seen in the corresponding Mendelian conditions. Secondly, in both diploid conditions, the decline levels off at a notably higher level than in either of the haploid conditions. (Haploid-Or: $M = .007$, $SD = .006$, ranging from .000 to .022; Diploid-Or: $M = .017$, $SD = .009$, from .002 to .030; Diploid-CoD: $M = .016$, $SD = .010$, from .000 to .040).

Tests conducted on the variance of trait values confirm both of these observations. Groups displayed significantly dissimilar rates of evolution, as indicated by strong interaction effects with time (Listen: $F(60) = 41.3$; Talk: $F(60) = 11.7$, p 's $< .001$). What level of variability does each method leave one with? Tests just looking at the end of the simulation, from time-step 2500 through 3000, showed significant differences between the three sorts of populations (Listen: $F(2) = 11.64$, $p < .001$; Talk: $F(2) = 6.23$, $p = .004$). Post hoc tests reveal that Haploid-Or populations had less variable trait values than in either sort of diploid population (for Diploid-Or, $p < .006$; for Diploid-CoD, $p < .05$), which did not differ significantly from each other ($p > .1$). See Table 1 for a summary of all the analyses conducted.

Another curious difference that appeared in the Haploid-CoD condition was that it supported fewer agents than any of the other groups (p 's < .001). On the 3000th time step, the average size of Haploid-CoD populations was 474.9, as opposed to 592.4, 588.5, and 587.2 for Haploid-Or, Diploid-Or, and Diploid-CoD, respectively. This gives the hint that Haploid-CoD populations, in being unable to evolve, were not doing as well as the others in a real sense. Figure 2.5 charts average population sizes over the course of the simulations.




Figure 2.5. Number of agents supported by the environment for each condition, over time.

Table 2.1. Summary of Inferential Statistics

	Talk		Listen	
	T1-T3000	T2500-3000	T1-T3000	T2500-T3000
Means				
Groups	F(2) = 12.45 p < .001	F(2) = 8.16 p < .001	F(2) = 49.73 p < .001	F(2) = 7.09 p < .002
Post hoc comparisons				
H-Or & D-Or	p = .041	p = .513	p < .001	p = .012
H-Or & D-CoD	p < .001	p = .001	p < .001	p = .004
D-Or & D-CoD	p = .061	p = .040	p = 1	p = 1
Time	F(30) = 170.50 p < .001	F(5) = 2.81 p = .018	F(30) = 886.70 p < .001	F(5) = 15.06 p < .001
Groups x Time	F(60) = 2.35 p < .001	F(10) = .43 p = .930	F(60) = 25.50 p < .001	F(10) = 2.12 p < .024
Variance				
Groups	F(2) = 29.76 p < .001	F(2) 6.23 p = .004	F(2) 94.23 p < .001	F(2) = 11.64 p < .001
Post hoc comparisons				
H-Or & D-Or	p < .001	p = .005	p < .001	p < .001
H-Or & D-CoD	p = .122	p = .045	p < .001	p = .025
D-Or & D-CoD	p < .001	p = 1	p < .001	p = .145
Time	F(30) = 223.63 p < .001	F(5) = 1.67 p = .142	F(30) = 784.08 p < .001	F(5) = 2.06 p = .072
Groups x Time	F(60) = 11.69 p < .001	F(10) = .36 p = .961	F(60) = 41.28 p < .001	F(10) = .64 p = .783

Table 2.1. Summary of inferential statistics. Results of all eight analyses of variance reported above are summarised here. Two tests were applied to both mean scores (Means) and variance of mean scores (Variance) for each dependant variable (Talk and Listen). One test considers the entire span of the simulations, from the first time-step to the 3000th (T1-T3000). The other considers only the last 500 time-steps, from the 2500th to the 3000th (T2500-T3000). Each test is a 2-factor ANOVA (3 groups x a number of time-steps, 31 or 6). F statistics, degrees of freedom (in parentheses), and probabilities of type 1 errors (p) are listed for each main effect (Groups, Time), and the interaction of these (Groups x Time), as well as for post hoc tests between groups (Bonferroni). Three post hoc comparisons are relevant in each case, one between each pair of groups (e.g. “H-Or & D-Or” is to abbreviate a post hoc comparison between Haploid-Or and Diploid-Or groups).

Applying Diploidy and Dominance

Since the influences of diploidy and Mendelian dominance promote evolution, they may benefit certain engineering applications that utilise evolutionary mechanisms. The advantage for finding solutions to problems that these processes are anticipated to confer is in insulating a search from getting stuck at some local optimum. Diploidy does this by preserving a reservoir of variation. The anti-conservative bias in Mendelian dominance will add further reach to the process.

The point that a diploid representation offers benefits in dynamic, time-dependent problem situations with moving optima has been aptly made with the work on diploidy in GAs (Singh, 2002), treated above. The other situation in which diploidy especially helps is when the problem is complex. It is not only moving optima that can trap evolution that moves too quickly, but those that take very much evolution to achieve. “A harder problem requires a larger diversity of structures in the population” (Shimodaira, 2000, p.49). Consider the Listen scores charted in Figure 2.2. None of the four sorts of populations reached what might be supposed to be the optimum score of 100%, yet at the end of the simulation, haploid populations had entirely converged.

In simpler problems, with fitness definitions that are unambiguous, unimodal, and unmoving, a haploid algorithm may outperform a diploid version (Yilmaz & Wu, 2002). Even the haploid-Mendelian version reported upon here achieved higher scores than the corresponding diploid version, which would most often be the goal (though not, for example, here!). In cases where they

perform equally, the haploid is to be preferred as it is somewhat cheaper in terms of both time and space.

The redundancy in the diploid genome means that it occupies twice as much memory as the haploid. Though it entails no significant hike in processing costs over the same period, if the intent is to bring it to convergence, the processing demands for the diploid representation will, naturally, be greater. However, it will anyway be of the same order of complexity as the haploid. Dominance, as described here, adds neither space nor time costs over codominance.

The practical benefits afforded by diploidy and dominance invite inventive bastardisations of these ideas. For instance, accepting that the benefit of complete dominance lies in its radical bias, try making the bias even more radical. Instead of “A or B”, try “A or B +/-10%”. Likewise, in the way diploidy maintains a store, this study invites the engineer to try triploidy or other polyploid strategies. This is not without precedent in nature. In some animals, the sexes even have dissimilar ploidy amounts, demonstrating how even fractions are possible, as well, such as the heritability of hymenoptera. Dasgupta and McGregor (1992) is such an example of an innovative response to diploidy in genetic algorithms.

Conclusions

Some distinct overall conclusions can be drawn from the pattern of results acquired, about how each of the four methods modelled allows evolution to proceed. Diploidy fulfils its promise of preserving variability. The haploid-Mendelian condition heads for high ground early in the simulation. It compromises variability for high scores. It is easy to imagine situations in which this would be a good strategy; however, it commits its possessors to the current niche. Diploidy, however, seems to make its possessors somewhat shy of the top. Trait values stabilised at a markedly lower level in the diploid conditions. In the face of the conflicting implicit goals of scoring highly, and preserving variability, haploidy chooses the former; diploidy, the latter.

Mendelian dominance, likewise, was seen to preserve variability. The difference is particularly striking nearer to the beginning of the simulations, in the first 1000 time steps (see Figures 2.3 and 2.4). Just as the haploid-codominant species loses its ability to improve by throwing away all its variability, the diploid-codominant species, too, squanders much of its initial variability early in the simulation, and ultimately fails to match the average trait value levels achieved in either Mendelian condition (see Figure 2.1). Much more gradual declines in variability are seen in the corresponding Mendelian conditions. Nor does this insulation of variability come with a compromise to adaptivity. Diploid-Mendelian populations scored at least as highly as the diploid-codominant, and significantly higher in Talk scores. So by both criteria, complete dominance appears to be an improvement over codominance.

The improvement that Mendelian dominance imparts over codominance to the haploid species is more extreme. Indeed, the haploid-codominant species must be faulty, for it forbids directional evolution. The lesson of the haploid-codominant agent seems to be that either of these ways of preserving variability, either dominance or diploidy, rescues it from stagnation. Hence, both the diploid-codominant and haploid-Mendelian conditions performed well.

Mutation was left out of our model. To the extent that diploidy preserves variability, and mutation creates it, the differences observed between the ultimate variability in the haploid genome versus that in the diploid would have been greater had mutation been included. Mutation would have given the diploid genome more diversity to hide; and the haploid, more to squander.

These results echo consistent findings in classical population genetics. Diploidy and dominance strongly affect evolution. If they did not, one would not be obliged to include them in a model. As they do, where accuracy is an issue, a modeller should consider these aspects of sex, particularly where studying evolution in sexual agents. Further, since these influences promote evolution, their addition to a model may benefit certain engineering applications.

CHAPTER THREE

THE COSTS OF COUPLING

Among the mysteries of sex is how it may account for its costs. A notable difference between sexual and asexual reproduction is that sex imposes an additional behavioural requirement. It obliges that two sorts of agent couple, that the two sexes be led to each other. Whatever the difficulties in mating, a sexual organism is faced, in the first instance, with the problem of finding of an appropriate mating partner. This may potentially be a very difficult problem. And asex has no such need. Since in nature, many sexual species compete directly with asexual species, nature must have some sure methods for making up for the apparent cost of coupling. This is the problem in focus: what behavioural requirements does the need to couple for reproduction impose on a species?

Coupler's Dilemma

Coupling involves a dilemma, the choice of whether to look for a mate or not. This is an important decision, for life is hard, and energies spent on chasing mates are taken from efforts toward the normal business of survival. Good foraging time and energy is forfeited whenever one invests in coupling. The dilemma is more complex than this, though, for coupling involves two interacting interests. A principal feature of this dilemma—abbreviated here as the Coupler's Dilemma—is that the two parties share the benefits of mating, but bear the costs individually. Both parents, after all, have the same interest in every fry produced. The question will arise of whether there is a way of reaping this reward while avoiding the costs. After all, if two things wish to find each other, it is enough for one of them to look.

That attraction can work as well when it is one-sided, as when it is reciprocal, is the guiding intuition behind this chapter. The interest is to identify aspects of the need to couple that can lead to predictions about the behaviour of a general case of a species so constrained. To so generalise requires a distinct, qualitative criterion, provided here by the expectation of a division of labour in coupling efforts. Under what circumstances will sex-based asymmetries in coupling emerge? This is the criterion question by which each addition to the model will be evaluated.

Like every other step on the road to reproduction, coupling is a cooperative venture, one that involves an interaction of motives and dependencies. In this way, the Coupler's Dilemma is analogous to the

cooperation games of Game Theory. Though no attempt will be made here to apply Game Theory to the Coupler's Dilemma, it is helpful to borrow the terminology of cooperation and defection, and with it, the hint of a deeper applicability. To *cooperate* in the Coupler's Dilemma means to actively seek a partner; to *defect* is to ignore the common goal of seeking sex, and concern oneself with one's own needs for survival. This comparison is appropriate, for it accentuates that it is a social dilemma in which dependent parties decide independently whether to contribute to a collective task or not, and that different combinations of actions lead to different outcomes for all. Moreover, phrasing the matter so borrows insight from that perspective, leading to the identification of at least one persuasive factor.

Sex Strike

One-sided sex may be fine for the population, but are your offspring going to be a part of that population? If one wastes no time on finding mates, he might yet be found by others, and seem to win the best of both, of finding a mate and reproducing, and at the other's expense. Yet complacency brings with it dependency on the other sex. A complacent agent has no guarantee that it will reproduce to capacity, or even at all. Any individual who actively seeks will miss fewer mating opportunities, and so have a selective advantage over non-seekers. As a disgruntled worker who refuses to work will be replaced, an uncooperative agent who refuses to seek sex will leave fewer offspring.

However, a coalition of disgruntled workers have power to enforce cooperation. So too might a coalition of agents strike for more cooperation in others. Who might strike in Sugarland, and against whom, if not one of the sexes against the other? If all Blues were to cooperatively *not* motivate sex, they would oblige the Reds to do so all the more, and so absolve themselves of that need.

However, as with any strike, it is foiled by the “scab” worker. A scab is “a worker who accepts employment or replaces a union worker during a strike” (Merriam-Webster, 2000). By cheating on the strike, one takes personal benefit, but weakens the effectiveness of the coalition. Given that each individual Blue might win a slight selective advantage by scabbing against his competitors, the other Blues, and pursuing sex, what is to keep them in line, and participating in a strike? This is again the problem of whether to cooperate with the common or tend to one’s own. However, it is a special case of this problem, with some features that favour coalescing: the two teams are clearly defined; and most importantly, what they are competing for is each other. This makes the essential difference.

Will the interests of the sex as a whole be able to outpace the selection pressures of the individual Blues? It depends on the strength of the competing selection pressures. However, it is not as unlikely as it at first may seem, if it is a matter of sex, simply because there are so many Reds about to be differentially affected by any change. When 100 Blues collectively still their urges to chase Reds, desirous Blues may win a slight selective advantage by missing fewer opportunities to mate, but it will be less than that thereby won by especially desirous Reds, for it is their market that is being cut into. Already there is a

difference of the sort predicted.

Strikes are only bound to succeed in situations that rely on heterosexuality, such as attraction. This is because such situations allow one sex to exploit the fact that the other is conversely affected by changes in within-sex competition in the other. Every relaxation of competition among Blues heightens competition among Reds, and vice versa. For instance, consider the arrival of a new Red to the population. This means tougher competition among Reds, and a lightening of the same in Blues. Likewise, increasing sexual desire in Reds heightens competition among Reds, while lightening it in Blues. This means that any agent that acts in the interest of its sex by lightening competition will find allies in the selfishness of all opposite sexed agents.

The more Reds that cooperate, the more Blues can get away with defecting, and so in turn, the greater the pressure for cooperation in Reds. Hence, any imbalance in investment here, gives one sex a ticket to defect. If Reds do not cooperate consistently enough, however, selection will favour cooperating Blues, and the union's designs will be foiled while Blues slip back into intermittent cooperation.

One reservation about this scenario might be that the other sex would find itself in the same position and act similarly, that between-sex competition would lead to a balance. This would be unlikely, however, for the pressure to defect changes with within-sex competition. The harsher the competition for mates, the greater becomes the pressure to pursue sex. And given that harsher for Reds means less harsh for Blues, any skew that appears here is expected to thrust the species into a diversifying positive feedback cycle, leading to only one stable

outcome: cooperation in one sex; defection in the other. See Figure 3.1 for a depiction of how these factors are expected to relate.




Figure 3.1. Depiction of the unbalancing relation between Desire scores and within-sex competition.

Note that this argument fails for most sorts of qualities. Such a strike could scarcely make one sex smarter or better hunters. There must be a discrepancy to exploit between the ways the sexes are affected in order for this sort of selection to operate. Any heightening of appetites for food, for instance, or for any resource all agents compete for equally, could only encumber the lot for the lot.

A “left holding the bag” argument could be imagined for harvesting, but see where it fails to muster a strike. Suppose that one sex were to erect a hunger strike, and demand that the other feed it. Human infants, for example, do this to

their parents and get away with it. Can one sex do the same to the other? The situation would again be that both sexes need the other to survive, and so those that feed their mates would excel, over those that do not. Yet, this situation invites defection—strikers feeding themselves—without the allegiance of selfish, opposite-sex rebels to back it, as seen in the case of attraction. In cooperating with the hunger strike, an agent is up against, not only the self-servitude of one's own sex, but that of the other, as well. As with the case of parental care, the hunger strike would require a commitment. One might even say of this hunger strike scenario that it sounds feasible. The point in raising it here was to note how much more feasible a sex strike is.

Though the sex strike will tend to lead towards asymmetry, it does rely on some slack in within-sex competition in one sex, and might yet be overshadowed where within-sex competition is high. Whether it does or not in a given circumstance is an empirical matter. Whether a sex-based asymmetry will emerge in coupling efforts is a primary question of the present work.

Analogy with Parental Care

An analogous situation is seen in another duty shared, in species that have it, by males and females: parental care. In many species, the young are born helpless. Something has to take care of them. The agent is faced with the dilemma of whether to care for its offspring, or to care for itself. The offspring must be cared for – natural selection will be emphatic about this. Selection might

be expected to strongly favour individuals that tend to their young. Yet, as both parties need this duty performed for them to have any progeny at all, if one parent were to withhold parental care, it would oblige the other to make up for it, at no cost to itself. One sex can do it, but can you trust that your partner will?

In many species, both sexes cooperate to raise the young; in many others, only one sex does. What is consistent here is that, where such a division of labour emerges, it is drawn on lines of sex. Being absent may be a feature of fatherhood or of motherhood, but rarely will a species have a mix of bad fathers and bad mothers. What seems to decide this matter is whether one sex is committed to staying with the egg or young longer. I repeat, if your opponent cannot risk a double defection, by defecting, you oblige him to cooperate. And if the opponent is sure to cooperate, you can defect with impunity. In mammals, for instance, where fertilization is internal, the female is “left holding the bag” of parental care. Hence, kittens are exclusively the toil of their mothers, the tom long gone. In fish species that care for their young, it is often the male that finds himself fertilising untended eggs, and so is stuck with the chore of protecting them. Meanwhile the mother has long abandoned. Consider, in contrast to these, birds of flight, in whom concerns of weight forbid the carriage of eggs. Neither the mother nor the father of a brood of young can risk leaving parental care to the other (possibly because it really requires two, but not necessarily), and so they are both bound to serve. The case of parental care is different than the coupling problem in at least one important aspect: something external to the dilemma commits one sex to cooperation.

No Dilemma

The dilemma is one that the Purple agent faces, but that many species would not. Several characteristics of the sexual lives of some animals would essentially nullify the cost of defecting. In Purples, reproductive success is a measure of quantity – the more offspring, the better. However, many species are not like this. In species that care for their young, as the most interesting animals do, there are limits to the number of offspring that can be well reared, producing beyond which can reduce inclusive fitness, rather than increase it. A human couple, for instance, had better not produce 20 children. In this case, fewer mating attempts could hardly hurt. The same is true of a species if a limit is imposed over some interval, such as where a mother will bear, at most, a single litter per season. In this case, it is of little worry if some mating attempts are passed over. Thirdly, monogamy too removes the force of this dilemma.

It may also be that one sex will be faced with the dilemma, and the other not. For instance, when one sex rears the young alone, its reproduction is better kept low, whereas the other may still benefit from quantity. This is a much talked about pattern, one that includes things like cats and dogs and many other mammals and fish. Note that this situation provides an extra push for the asymmetrical solution to the Coupler's Dilemma, because the one continues to gain with each extra effort. This implies an interesting dependency between the outcomes of the Coupler's Dilemma, and the analogous dilemma of whether to invest parental care. When renegeing on parental care, one sex may be consigning itself to the toils of coupling.

Method and the Purple Agent

The JAWAS modelling platform is used to simulate the evolution of a progression of heritable aggregative and segregative behavioural influences, some that allow the components of attraction and repulsion to be dissociated. These were designed to probe the very simplest hypothetical elements of sexual attraction. Each behaviour is allowed to evolve independently in each sex, and the extent that each does is taken to indicate the way a behaving organism is likely to evolve under the sway of such circumstances. Conditions are identified under which a sex-based asymmetry in sexual motivation emerges.

The star of this chapter is the Purple agent. The Purple was designed to be *merely* sexual. This means that the rest of it is to be arbitrary, and as general as possible. An arbitrary agent has many arbitrary details, such as how far it can sense, or how long it lives, any of which may well affect outcomes. There can be no pretence of being exhaustive. Whatever parameters the system allows or does not allow to be varied, having varied them all, one could scarcely claim to have been thorough. One can yet be systematic with a few variables of interest, and licentious with the rest, picking parameter settings that especially pull out the effect to be demonstrated, not forgetting that the intent was nothing other.

Life to the Purple Agent

At the top of every time step, each agent decides whether to forage or to pursue some potential mate. If he will forage, he looks about him, to the extent he can sense (10 cells in each of the four cardinal directions), and picks the cell with the most resource. He then goes there, and eats what he can of it. If instead he chooses sex, he again looks about him, but is now judging the appeal of the other agents. He picks the most appealing agent in sight, goes to its cell, and attempts to mate. Appeal is a matter of several traits described in coming sections, and agents may disagree about what is appealing and what not. Mating is successful if the agent is of the opposite sex, and if neither has yet reproduced this round. Each agent is allowed to produce at most one offspring in a given round. If a success, a new agent is born and added to the population, mature and ready to make adult decisions in the coming round. It inherits all genetic traits from its two parents. Agents die if they fail to harvest enough sugar, or upon their 200th birthday (i.e. the maximum life length is set at 200 time steps). All this is happening simultaneously in several hundred agents in a small grid world.

There are several things missing from this picture. The first omission to note is any of the concerns of mating – beyond the very first, that of finding the mate. Success in finding a mate leads straight to success in mating. The other agent has no recourse to reject the union, demand that it bring it food or dance a jig, or any of the other aspects of mating in natural organisms. These omissions are intentional, for the model is about how agents are led to appropriate mates, not what they do with them once found.

Genetics to the Purple Agent

The Purple agent is equipped with 102 genes, any of which can be set constant, varied, or disabled entirely. This flexibility allows a great number of experimental situations, several of which are reported upon below. As the developments are cumulative, they are discussed as they appear in the course of the problem. All genetic matters involve attracting and being attracted to potential mates. When it comes to harvesting resources, the criterion for survival, all agents perform identically—all have the same metabolic rates, consumption requirements, speed of movement, efficiency of sensors, and so forth—and that part of their codes will not be permitted to vary. This means that evolution will never make them better harvesters. This carries some further repercussions on the evolution of attraction, discussed below.

Purple agent possesses the diploid-Mendelian genome from Chapter Two. Unlike the agents of Chapter Two, traits in Purple agents are subjected to mutation. New Purples entering the population will inherit their parents' genes modified by the addition of a random value taken from a Gaussian distribution, with a standard deviation of the mutation rate, here set at 0.05. As it is a diploid genome, each agent carries two alleles for each trait, for a total of 204 values remembered. Some of these are real numbers, as those in SugarAgent; others are Boolean values. Boolean traits control presence or absence of a quality; real-valued traits represent the quantity or strength of an influence.

Dominance and Recessiveness

With the introduction of Boolean genes, which code for presence and absence of a trait, dominance relations that allow for consistently recessive or dominant traits become easier to preserve an accurate notion of. Recall that the agents' genome is a double list, with two values for each trait (this double spot on the data structure is a "locus", in the terminology of genetics). For the quantitative traits, one dominating means just to consider one of them, and ignore the other. In the Boolean traits however, a trait can be recessive. This means that it is expressed only when both elements at the locus are filled by copies of the same allele. Likewise, a trait can be dominant, meaning that only one such allele needs to be present for the quality to be expressed. For the course of this work, the influences of this aspect of allele dominance are not explicitly investigated. It does leave hints for future work, though.

Population Density

Where population concentrations (i.e. agent : space ratio) are high, many successful Red+Blue couplings will occur incidentally, and aggregative behaviour becomes superfluous. Indeed, in many simulation environments, including VUScape, there are no rules for explicitly bringing agents together for mating, relying instead on high population densities. The simulations conducted here are under one of two population densities, neither of which supports populations

without some form of aggregative behaviour. The high density situation is at an agent to cell ratio of 1 to 8 (1250 agents in a 100 x 100 cell grid); the low density ratio is 1 agent to 125 cells (320 agents in a 200 x 200 cell grid). These densities are thought to be dissimilar enough to represent this dimension.

Genetic Drift

Genetic drift, also known as the Sewall Wright effect, refers to the part of evolution that occurs outside of Natural Selection, due to cumulative sampling error (Thain & Hickman, 1994). One's past determines one's future, and particularly if one is an evolutionary algorithm. Random sampling errors can magnify over generations to produce outcomes sparked merely by chance. Recall from Chapter Two that, in the diploid genome, at least half of the gene pool is anyway wholly invisible to natural selection. This makes it all the more susceptible to effects of cumulative sampling errors.

When a genome hosts large numbers of traits, such as the seemingly paltry 102 of Purple agent, genetic drift begins to exert a sizeable influence upon evolution. There is too much to optimise simultaneously. The chances of coming to optimal levels on more than a few traits are slight, especially with low population sizes.

This is both a problem to be guarded against, and part of the answer to my question. It is part of the answer because it permits many transient deviations from ideally rational solutions. It assures that any pair of genomes will be

lopsided at least transiently, whether or not it thereafter converges upon a balanced solution. As the bilateral cooperation solution to the Coupler's Dilemma problem discussed above appears to be a viable, possibly stable, alternative to either asymmetrical solution, it is helpful to be assured that all will get their chance to shine or fade.

Further, genetic drift will reveal itself to be a powerful determinant of how agent populations evolve, in unexpectedly systematic ways. Specifically, it can push toward divergence between the sexes. How it does so is discussed later.

The problem with genetic drift is that it conceals weak effects, or rather, swamps them. These sways can be curbed by allowing only a small amount of the genome to be varied at a time, controlling other factors.

Desire and Other Dependent Measures

Desire

The first issue to confront is the cost of coupling. Recall of the SugarAgent that it must harvest in order to survive, and mate in order to reproduce, but that it can do both in any round. In Purples, these functions are in conflict, such that only one of the two can be enacted in any given round, hence the agent's dilemma. If an agent spends the round in sex, it may find a partner, and reproduce. If it sticks to work, it will feed its grumbling belly. Where nutrition is scarce, choosing to move on any other basis can be a hindrance to an agent's

survival. Such is the cost of coupling. Yet, reproduction is the only route available for gene survival.

An agent's Desire score is the condition for all its coupling behaviour. Is the agent in the mood? If it is, it will forego food, in favour of seeking sex. If not, it will ignore other agents and forage. Desire, then, is a measure of the agent's response to the Coupler's Dilemma. As with other traits, each agent's response to the dilemma is represented by a real number. This is a measure of how often the agent will choose to spend the round seeking mates, and how often it will forage.

At the beginning of each simulation, the Desire value in all initial agents is set at a standard 0.5. This means that the first generation of agents will spend half of their time seeking mates, and half their time foraging. With every new birth, a little variability will be added to the gene pool through mutation upon parents' traits, some to be mixed into the coming generations, some to be weeded out by natural selection.

In the simulations reported upon here, life conditions are harsh, so subsequent generations will evolve to spend more of their time foraging. This makes the point that the cost of coupling is dependent on other factors. However, when the limits of harsh living are approached, all relativity flies – one must eat after all. Nor are Purples exempt of this. So Sugarland was stocked to keep Purples on the edge, such that every mating attempt has real cost to them. “On the edge” was achieved in the simulations by allowing populations to reach their own, natural population ceiling, in which there is a cruel equilibrium of births and deaths. This solution has the dual benefits of keeping the cost of coupling both high and standard. It further allows the cost to be manipulated by

progressively imposing an artificial population ceiling, giving agents a bit of a break.

We are interested in the case where the sexes will diverge on Desire scores. More specifically, the interest is to identify factors that will push the Desire score in one sex down to the point where it never seeks, and it is only by the other sex's efforts that mates are found. The motivation for requiring the one sex to entirely give up on coupling is with the hope of identifying *qualitative* effects on behaviour. Quantitative differences might also be interesting, but considering the distance between the Purple and what it hopes to model, attempts to draw generalisations, however tentative, on anything but qualitative criteria would be tenuous.

Later developments require that Desire judgements be determined for each agent at the beginning of the round, before any are permitted to act. This means that the execution cycle of Purple agent follows a third subcycle, beyond the two of SugarAgent.

A Note On Measurement

As the Boolean traits indicate whether a prescribed action will be performed or not with its two values, a "yes" and a "no", many of the traits represented by a real number also indicate whether an action is performed or not. The difference is that real numbers allow differentiation between gradations of likelihood that the action will be performed, as in SugarAgent. These traits are

used to represent an agent's propensity to act in the prescribed way. For instance, an agent might act so only half of the time, or on 30% of rounds.

However, traits are unbounded and so may evolve to be any value. This means that only an interval of trait values make a difference in behaviour. The nice feature of this is that it allows traits to settle upon any of the entire range of probabilities (to two decimal points), or to follow a gradual path to performing the action either always or never. For convenience, this interval is between 0 and 1, such that a trait value of 0.78, for instance, indicates that the agent will engage in the action 78% of the time. Any trait value above this interval indicates that the behaviour will be enacted with every opportunity; values below indicate that the agent will never perform the action.

Why not cut trait values at 0 and 1? The main motive for letting mutation do with these values what it will is that we are interested in the possibility of a committed zero, and unambiguously differentiating this from just very low positive values. We want to allow "no action" to be a feasible alternative.

Imposing a floor at zero forbids this, leaving us to wonder whether the optimal value for the agent was really just very low, or whether zero would have been selected had it been allowed. You say, though, that imposing a floor *does* allow zero. Not really, for mutation begins to behave differently as it approaches imposed boundaries. Recall from the definition of mutation given above that it is the adding of a value taken randomly from a Gaussian distribution. Consider what a boundary would do to such a distribution, and how it would distort what values are passed on. Specifically, it biases values to be away from the boundaries, and not in a small way. Even without mutation, such boundaries

would affect distributions artificially with what are known as ceiling effects and floor effects. These say the same thing: as boundaries are approached, variance drops artificially. With mutation, however, this situation is exacerbated, as it pushes values especially away from the boundary, again artificially. I can think of no motive for restraining trait values so.

Leaving a trait unrestrained makes it a more sensitive measure. Given that these traits are the dependent variables, sensitivity of their measurement is especially important.

The Traits

Heterosexual attraction is something that can be characterised functionally, and with some hope of doing so exhaustively. The traits of the Purple agent are designed with such an ambition. Each Purple Agent possesses several traits that need to be abbreviated, as well as some conglomerations of traits that encompass the action of many genes. These are given in Table 3.1, and explained in subsequent sections.

Table 3.1. The Traits of the Purple Agent

<u>Trait group</u>	<u>Traits included</u>
Desire	Desire, DesireBlue, DesireRed
Approach	Approach, ApproachBlue, ApproachRed
Avoid	Avoid, AvoidBlue, AvoidRed
Attract	Attract, AttractBlue, AttractRed
Repulse	Repulse, RepulseBlue, RepulseRed

<u>Beauty</u>	attract1Up attract2Up attract3Up attract4Up attract5Up attract6Up	attract1Bp attract2Bp attract3Bp attract4Bp attract5Bp attract6Bp	attract1Rp attract2Rp attract3Rp attract4Rp attract5Rp attract6Rp	attract1UQ attract2UQ attract3UQ attract4UQ attract5UQ attract6UQ	attract1BQ attract2BQ attract3BQ attract4BQ attract5BQ attract6BQ	attract1RQ attract2RQ attract3RQ attract4RQ attract5RQ attract6RQ
<u>Ugliness</u>	repulse1Up repulse2Up repulse3Up repulse4Up repulse5Up repulse6Up	repulse1Bp repulse2Bp repulse3Bp repulse4Bp repulse5Bp repulse6Bp	repulse1Rp repulse2Rp repulse3Rp repulse4Rp repulse5Rp repulse6Rp	repulse1UQ repulse2UQ repulse3UQ repulse4UQ repulse5UQ repulse6UQ	repulse1BQ repulse2BQ repulse3BQ repulse4BQ repulse5BQ repulse6BQ	repulse1RQ repulse2RQ repulse3RQ repulse4RQ repulse5RQ repulse6RQ

Table 3.1. The traits, grouped in behaviours, each referring to a subsequent section, where the trait is explained. In Beauty and Ugliness, traits come in two broad sorts. Boolean values control the presence or absence of the behaviour or influence, indicated by a “p” in the gene’s name. A “Q” in the same spot indicates that the trait controls the amount or strength of the effect. “B” in the trait name indicates that it expresses exclusively in Blues (e.g. Attract6Bp means p Attraction to marker 6 by Blues). Likewise “R” indicates expression in Reds, and “U” indicates that the effect is sex-universal (e.g. repulse5UQ refers to a Q attraction to marker 5 that manifests in both sexes).

Sex-specificity

As the beginning question is whether sex-specificity will emerge in coupling behaviour, all of the behaviours studied are allowed to appear in each sex independently. This means that each of these behaviours is controlled by

three genes, corresponding to each of the three possible, functionally distinct, sex dependencies: some genes affect only Blues, and are carried dormant in Reds; some affect only Reds; and some are expressed in both sexes. For the tests reported upon below, only the first two of these sorts of genes, those specific to one sex, are used. The sex-general sort was added primarily to test for spurious divergence, but that test is not treated here.

Simple Aggregation and Segregation

Beyond the few traits that are used in the demonstrations recounted here, several additional traits are a part of the agents' repertoire of potential behaviours. Simulations with these are not treated here, for later additions make them obsolete. However, as they are part of the Purple agent, and as that it is a part of this study, they are briefly described here.

Simple model: Approach and Avoidance

The simplest behaviour for promoting contact among the agents is to have them attracted to each other indiscriminately. Approach is an agent's bare, indiscriminate aggregative behaviour. An agent's Approach score indicates the likelihood that it will use its movement on a given turn to especially enter a cell occupied by another agent irrespective of what sort of agent it is. Avoidance is an

agent's tendency to actively avoid other agents. The Avoidance score indicates the likelihood that the agent will use its movement on a given turn to especially enter a cell that is occupied by no other agent.

Naïve model: Attraction and Repulsion

Attract is an agent's tendency for heterosexual aggregation. An agent's Attract score is the likelihood that it will use its movement on a given turn to especially enter a cell occupied by another opposite-sex agent. Repulse is an agent's tendency for homosexual avoidance. An agent's Repulse score is the likelihood that it will use its movement on a given turn to especially avoid a cell occupied by another same-sex agent.

If one were modelling something further along in sexual attraction, one could for most purposes be content with Attract and Repulse as a model of heterosexuality. However, differences that emerge between this and the extended model it abbreviates, of the next section, attest to the notion that there is something of interest along the way to heterosexuality, something not seen when heterosexuality comes pre-assembled, as in Attract and Repulse. Apart from exposing this difference, having this abbreviation of the extended model affords a partial solution to high genetic drift for the sake of some tests. It allows, for instance, all interesting combinations of the above five traits to be exhaustively tested, forbidden by the complexity of the extended model.

Beauty and Ugliness

For the sake of many purposes, *Attract*, given in the previous section, might be an apt model for heterosexual attraction. However, it fails in the present case, as the interest is to identify an influence hiding somewhere within it. Attraction can be further analysed into its functional elements. This is what this section attempts to do.

The first dissociation to make in this is between the subject and object of attraction. Wherever attraction is innate, there are two things that might evolve in tandem: the thing that attracts, and the thing that is attracted. In *Attract*, the former is taken for granted. In *Beauty*, both the mark of beauty, and recognition of it are permitted to evolve independently of each other. A marker may become an attractant, or a repellent, or both. A marker is an attractant if agents evolve to be attracted to it; a repellent if agents evolve to be repulsed by it.

Beauty is in the eye of the beholder, as well as on the surface of the beheld. The point that it is not merely in the beholder has been made for particular species, including fireflies (Lloyd, 1981) and humans (Langlois, et al., 2000). However, for the present purpose, where the target is not a particular species, but a general case, any such demonstration is neither sufficient, nor even necessary, for one can at least be sure that wherever beauty may actually lie, it is nowhere *other* than in these two places—beholder and beheld—and most likely it lies with both. The model, then, is obliged to allow attractiveness to evolve in both subjects and objects, and in nothing else.

Purple agents are attracted or repulsed not by whole agents—they are no longer magically aware of which is male and which female—but by various markers that appear on some agents and not others. These markers are modelled after the “releasing stimuli” of classical ethology (Tinbergen, 1965). The notion of the releasing stimulus is that some sensory stimuli directly elicit, or “release”, a specific behavioural response that the animal is evolved to perform under such circumstances. One example might be that the shape of an egg elicits nesting behaviour in a species of bird. The same stimulus might also release predatory behaviour in certain species of snake. Releasing stimuli are especially convenient for attraction, for the stimuli that must release sexual behaviour is not some objective, external thing in the world, but in the opposite sex, which is within control of the genes that make up the species. This makes the matter more complicated, for it relies upon a dynamic relation between the sexes. Even so, this dynamic frees the matter somewhat from its details, such that the relation is similar across species.

The Implementation

To fully represent the beholder and the beheld independently, and to moreover allow sex-specificity, requires a quickly increasing number of genes to work independently of each other. No longer is the attractive or repulsive thing a whole agent, but a marker that may appear on some agents in the population, and not others. Each marker can, as the traits mentioned above, be specific to

either sex, or be general to both. Thus, three distinct markers are required. As the Purple will need another three markers for the simulations of the next section, this number is doubled. These other markers are covered in the next section, where they arise. This section deals with the first three. Of each trio of markers, one is sex-general (U-Mark), one appears only in Blues that carry the gene, and is carried dormant in Reds (B-Mark), and the last is expressed only in Reds (R-Mark). All markers are recessive traits, meaning that they express only when two marker alleles are present at the locus simultaneously.

To each of these six markers, there are three potential sorts of attraction, and three potential sorts of repulsion, again, one for each of the three functionally distinct possibilities of sex-dependence. This results in 18 attractions to the various markers, and 18 repulsions. These traits are Boolean values that control the presence or absence of the attraction or repulsion. Each of these attractions and repulsions are associated with a second, qualifying real value that indicates the strength of the effect. These are allowed to evolve independently, for another 36 genes. These 72 response genes are given with their identifying names in Table 3.1.

To clarify, each attraction and each repulsion is represented by two genes: one that controls its presence or absence (p-genes), and a second that represents its strength (Q-genes). Looking at these genes in the population, the former, a Boolean value, is measured in terms of its frequency; the latter, a Real number, in terms of average quantity. In the simulations reported upon here, p-genes are expressed in about 50% of initial agents, except where stated otherwise. These traits are dominant, meaning that they are expressed whenever one such allele is

present. Q-genes are subjected to mutation, but p-genes are not. The rationale for this is that the expression of the trait is a product of both genes, so it is enough to have mutation in the one. Mutation for the Boolean traits is, moreover, more disruptive, for it can only change between extremes, there being only two possible values.

At the beginning of each round, what each agent will do is determined by its Desire score. If it will forage, there is nothing to do with Beauty or Ugliness. If it will seek mating opportunities, it makes its decision about which cell to move to on the basis of the markers on agents near enough to sense. It can see 40 cells from any spot in VUScape. An agent guided by Beauty will isolate the agent that it finds most appealing, and move there. What it finds appealing depends on what markers are expressed in the agents about, which of these markers the agent fancies, and how much. The effects of the various markers are additive, so an agent expressing two markers will be especially attractive to agents that possess both attractions. In the Purple, any or all of the traits can be run simultaneously. In the simulations reported upon here, at most only Ugliness and Beauty are active together. These combined lead to avoiding the ugly, and approaching the pretty. Mostly these two will not be in conflict. Conflicts are resolved by comparing the strengths of the conflicting traits. For Beauty and Ugliness, if the ugliest agent in sight is uglier than the prettiest is pretty, and these happen to be on the same cell, the actor will forego that cell for another. Otherwise, it will brave the beast for the beauty.

Experimenting with Beauty

With the addition of Beauty and Ugliness, the flexibility of both the Purple and VUScape combine to allow hundreds of distinct simulations. Some general findings from these are reported upon here. These will be kept general where possible, only narrowing on particular simulations to note sensitive details.

Result One: Everyone is Beautiful

The first thing to notice about indicators of attraction is that they spread very readily, even greedily. When allowed to evolve unconstrained, they quickly come to spread unto ubiquity, even from very low starting frequencies – even from zero, where mutation is enabled. Very shortly every Blue has the B-Mark, every Red has the R-Mark, and everyone has the U-Mark. Likewise, reactions to each marker soar. Attractions develop toward opposite-sex markers; repulsions develop toward same-sex markers. However, under these conditions, B-Mark perfectly indicates Blueness; R-Mark perfectly indicates Redness. When this happens, U-Mark loses its utility, and reactions to it slowly wane, until all that remains is essentially the naïve Attract and Repulse strategies, as defined in the previous section. This is a very robust effect, and happens seemingly irrespective of initial settings, so long as possession of markers is heritable.

A second interesting point about the evolution of Beauty is that the haste with which markers dominate the population is irrespective of whether they

benefit or hinder their possessors. Consider, for instance, the waning of interest in the sex-universal marker. Despite the fact that U-Mark is not a good indicator of a mate, it proliferates nearly as speedily as those that are. It does so because its possessors attract callers, half of whom happen to be opposite-sex agents. Even while attraction to an obsolete marker dwindles, until it is entirely extinct, the marker will continue to exploit that attraction. A look at the charts of the evolution of a marker, and the attraction to it, shows the latter to approximate the first derivative of the former.

Once Marks R- and B- have reached ubiquity, and U-Mark wanes, the effect of Beauty is to make all Reds attractive to Blues, and all Blues attractive to Reds. That is, it has settled upon the same rule as “Attract”, in the previous section. Likewise, the effect of Ugliness has become nothing more than “Repulse”. As such, populations of agents in this situation respond to the Coupler’s Dilemma as did agents under Attract and Repulse. Attraction is shared by the two sexes, each attracted to the other. This validates my model, to an extent, but poses a problem given that the intent was to pull the influences of the subjecthood and objecthood of beauty apart from each other.

Result Two: When beauty is held at bay

The markers had to be held back artificially to keep them variable. The solution was to remove heritability from the presence or absence of each marker. For the simulations reported upon in the rest of this work, the presence of a

marker in the agent is not a genetic matter, but is determined randomly, at a constant frequency. For most simulations, this frequency was set at a default value of 0.5, meaning that half the agents would possess the gene for the marker. Simulations run with markers set at any positive frequency lead to results similar to where they are heritable. High attractions develop to the opposite-sex marker; a milder, transient attraction develops toward U-Mark; same-sex markers become repellents.

A Lovelier Purple

Keeping beauty at a static frequency, irrespective of the dictates of evolution, is not satisfying, for the hope was to vary attractiveness and attractedness separately from each other, in the suspicion that there is something interesting hidden in the interrelation between them.

In addition to mere recognition of a marker as attractive, a thing can be more or less attractive. As with the other quantitative traits, the degree to which each marker is held can be represented and varied. In the simulations that follow, there is a real numbered heritable trait, represented by a Q-gene, corresponding to each marker indicating its degree. This trait is a variable, heritable coefficient that modifies reactions to a marker. So instead of all instances of a marker being treated equally, markers are allowed to grow to differentially exploit an attraction to it.

In the analogy with animals, this is the longer peacock tail, the redder gull beak, or the fuller moose rack. Such markers are sex-specific, and if they indeed release sexuality in opposite-sex conspecifics, they are like the sex-specific attractants in the Purple agent, R-Marks and B-Marks. Do such markers in animals become more attractive when exaggerated? Such excesses as peacocks' tails would be difficult to explain otherwise. More direct evidence for preferences for indicators of the extremes of masculinity or femininity comes from work on human facial attractiveness, which show consistent, intriguing patterns in the way femininity and masculinity affect attractiveness in humans (Perrett, et al., 1998; Perrett, May & Yoshikawa, 1994; Penton-Voak, et al., 1999).

Result Three: Runaway Beauty

Returning heritability to Beauty in simulations results in very striking, unambiguous trends. There is again an insatiable increase in Beauty, but unlike the mere presence of a marker in the simulations above, the degree of a marker has no strict limits, and so continues to climb, at a steady rate, throughout the simulation. So if, for example, the simulation is aborted at 3000 time steps, the degree of markers will be much less than if the simulation continues for 100,000 time steps. Even then, however, it is not finished evolving – the rate of increase is the same. This is a clear example of sexual selection. See section 3.7 for a discussion of sexual selection and the intuitions this study contributes.

Agents in these simulations are dividing their energies between foraging and seeking mates. Both sexes contribute to coupling, for a symmetrical solution to the Coupler's Dilemma.

Result Four: Asymmetry

Being able to manipulate frequencies of the various markers allows a stacking of the deck. When traits that indicate beauty are rarer in one sex than the other, an asymmetry emerges.

Consider a simulation in which the sex-general marker appears in 50% of the agents; the Blues' marker appears in 75% of Blues (B-Mark); the Reds' marker (R-Mark) appears in 25% of Reds. Where agents have access to a good indicator of appropriate mates, such as a high-frequency sex-specific marker, they can specialise their attractions to that marker. Selection is relatively slow here on account of the strength of the selection for the attractants, so where one marker directs agents reliably, attractions to others fall behind. Hence, Reds in this situation become highly attracted to the B-Mark, and mostly indifferent to the U-Mark.

To Blues, however, the situation is much different. They are under an even stronger pull by sexual selection driving B-Mark up, simply because Reds are able to find them more often. R-Mark is not a very reliable marker at 25%, and so Blues are often led to rely on U-Mark, which as often as not leads them to frustrated mating attempts. Reds are more successful than Blues at finding

appropriate mates because of the reliability of the markers. As every mating attempt carries the same cost, frequent failures drive Desire scores back in Blues. Meanwhile, the successes of the Reds' mating efforts lead to much higher Desire scores. This division grows until only one sex is motivating coupling.

So the Blues have got the Reds trapped! They win the game of attraction, and go on to reap the sugar of Sugarland, while Reds bear the whole burden of coupling. Any rare Red that tries to challenge the Blues' monopoly over sexuality will fail to produce a progeny. Only by being more aggressive than competing Reds can its ways be remembered in subsequent generations. The Red is doomed to cycles of ever tightening slavery to aggressively seeking sex.

Observations about Purple Beauty

Active passivity

The skew imposed in the frequencies of indicators in the previous simulation does not lead to the asymmetry where Beauty is held constant. In such simulations, even with the Blues' poor mating rates, they do not give up trying. The lesson here is that selection will not slacken their competition. That Blues seem to give up is illusory. What happened, rather, was that there were better grounds for competing for mating opportunities: becoming more beautiful. Beauty was anyway the stronger force, but in that simulation, it wholly overpowered the pressure to seek, and the pressure to forage pushed Desire

scores down.

Finding and being found are both effective methods in coupling. In no simulation in this study did selection forgive either sex of one of these grounds for competition. This leads one to expect that the peahen, for instance, could scarcely be unattractive to peacocks, and indifferent to sex, as she is sometimes characterised. The literature on sexual selection is used to talking about the matter as if sex were all in the male, that he exclusively were both the bothered and the beautiful. My results however suggest that he could not be, that you could not convince the peahen or doe to not get involved in this. The moment one does, she immediately wins a selective advantage over her passive neighbours. If she seems to be passive, it must be that there is competition among them to be the more beautiful, however drab they may appear to human sensibilities. Beauty provides a currency for sexual passivity.

Rampant Beauty

There is a subtle asymmetry between the way the objects and subjects of beauty evolve. The releaser, the marker, is parasitic on features of other agents, such that it grows at a rate commensurate with the frequency of attractions to it. Though the natural world would eventually curb such rampant selection, Sugarland does not.

One interesting consequence of this observation is that it changes the nature of the problem. When beauty loses sight of its starting point, it becomes

relative. Zero attraction has real meaning, and however long a simulation is run, the other traits—for instance, the subjects' attractions—all stay within grasp of it. The subject of attraction stays grounded. Beauty, however, never settles upon an asymptote, but keeps growing at a constant rate. Once a marker's magnitude grows to the point where the range of competing values are far removed from its zero, this zero ceases to be meaningful, and the trait loses connection with its ground. What difference might its being relative make?

It speaks against Good Genes explanations for beauty – that features are attractive because they signal genetic well being. Without a meaningful zero, beauty can scarcely be a measure of something else, such as health or genetic soundness, simply because it could not help exceeding what it was meant to be an indicator of. It would, in a short time, lead to every gradation of the trait indicating top health. So though possession of the trait might indicate health, it could at best indicate a presence/absence distinction. This means that selection has lost any grounds for discriminating between gradations. This issue is discussed in the later section on Sexual Selection, but to clarify the qualm about relativity, an analogous argument is given here as well.

One attractive quality that may help to illustrate the way relativity may disrupt an attractant's ability to give better information is the giving of gifts. The giving of gifts is attractive, and has won many a heart, whatever protestations to the contrary. And the giving of bigger gifts is even more attractive. This situation is fine for the caveman, where wealth has limits, and a wealthy youth, able to give the most expensive gifts, would most often have been an especially able hunter. But money has removed from real wealth its foundation. Wealth is no

longer just a matter of being well-fed, but of being wealthier than one's neighbours. Is the biggest prize going to come from the ablest provider? It might, indeed, but this will no longer necessarily be the suitor with the keenest eyesight or surest aim. It has become a social matter whether bigger gifts mean surer provision, or not. It is no longer that only the richest can be generous. Nor is the gift bound to handicap the contemporary giver, especially. To the caveman, the gift was expensive; to us, it is a trifle. We can all afford to give. In any case, the giving and appreciation of gifts remains a part of our psychology, however removed from its hypothesised function it now may be – and however removed from its zero. The difference is that now it is hard to say how big a gift needs to be, such that only shockingly grand gifts can say more than the mere act of giving. It is now the thought that counts, if only because we have lost count (of absolute values).

This example was only meant to be illustrative. The point is that relativity in an attractant foils its power to give more than nominal-level information about its wearer.

Obsolete markers

Even the tiniest attraction to a heritable marker rapidly spurs it into currency in the population. This is true even where both the marker and the attraction are very rare, and often the attraction will fade out of the population, particularly when it is to a marker that becomes obsolete. For instance, when a

marker exists in all agents, and another exists in only Reds, there may be no pressure for an attraction to the former to develop in Blues. However, this will not stop the marker from exploiting the attraction before it peters. This means that there will often be an accumulation of obsolete markers in a species, some of which may later go on to assume functions irrespective of their beginnings.

Epistasis

One reservation about the observation of flyaway beauty is epistatic relations between genes. In Purple agent, genes are assumed to evolve independent of all other genes. In carbon-based organisms, however, genes are linked to each other by bonds of varying strengths. Will beauty still fly off so violently if it will take other, linked traits with it? This is where beauty is expected to become complex, beyond what this sort of modelling could keep pace with. However, it is worth thinking about what this would do, how beauty could affect a genome. Be sure that many traits would run with it. From the apparent strength with which attractants evolve, this seems certain. Linkage with a beauty marker would most often lead a trait to a constrained range of values, to ingrain whatever happened to correlate with the marker. Weak effects of linked genes would be overshadowed; perhaps even some stronger effects would be, as well. By physical, visible traits, linkage with beauty would tend to lead a population to more resemble each other. This is consistent with the finding that averageness in human faces contributes to attractiveness (Langlois & Roggman, 1990; Rhodes & Tremewan, 1996). Likewise with behavioural, psychological traits.

The Short Story of Beauty

The short story of beauty is that any attraction in one sex will power the spread of its object in the other with a strength to overshadow virtually any other selection pressure, including that of the attraction itself. The rub is that this runaway attractant is reliant on the attraction to it. And being so much weaker an effect, it is in danger of being left low by the random winds of genetic drift. This makes the tall towers of runaway beauty unstable. Considering further that it may often be not in the subject's interests to be attracted to an attractant, many a tower may tumble, and the marker become a fossil.

Such an event, the felling of an attractant's pull, might seem of little import to the individual, at least to the subject. However, each such event will be a blow to the balance between the sexes, and could set a species upon a course to the asymmetrical outcome in the Coupler's Dilemma.

Smiles and Cool looks

Displays of commitment

There is a line of work in which simulations are conducted that plot different strategies for playing Game theoretic games against each other, in a “survival of the fittest” spirit, to see which strategies would survive to become stable (Axelrod, 1984). A memorable result from these studies was that a conditional strategy—cooperate with cooperators; defect on defectors—does remarkably well in the iterated Prisoner’s Dilemma. So might a conditional strategy do well in the Coupler’s.

The markers of this section allow agents to differentiate Desirous agents from un-Desirous ones. A Desirous agent is, of course, one who is about to cooperate in the current round; an un-Desirous agent is one who will defect. This allows the agent to take a conditional strategy towards coupling.

Note that this is different from the Prisoner’s Dilemma, for here agents see other agents’ intent, whereas the Prisoner’s Dilemma situation has the two players strictly ignorant of each others’ move until the chips are collected. Axelrod’s (1984) “tit-for-tat” strategy was conditional on opponents’ previous move. However here, albeit with some margin of error, the agent can see what other agents are going to do in the *current* round – which agents are about to cooperate and which defect. Note that the agent, too, has already made its own decision by the time this knowledge comes, but before it acts upon its decision. Lest this sound confusing, recall that all agents make their decision of whether to

act or not at the beginning of the round. This sort of marker just makes the decision transparent, such that when it comes time to act—eating time for the un-Desirous; mate-finding time for the Desirous—the agent’s choice is transparent to other agents.

Think of the game of Chicken. If you could be sure of your opponent’s decision, you would know what to do. Suppose you looked up to see your opponent, still speeding towards you—driving a convertible, let us say—not sitting tensely at the wheel, but standing upright on the seat, arms spread, apparently with the wheel and gas pedal rigged. She is displaying her commitment to stick to course. You would do well to yield here. Or if the opponent, instead, pulls over to the side of the road, you may drive on with impunity. Such indicators conditional on one’s choice change the game of Chicken. What will they do to the Coupler’s Dilemma?

Smiles and other indicators of cooperation

A smiling person is usually more attractive than the same person would seem without the smile (Cunningham, 1986; O’Doherty, et al., 2003). He might be smiling for any number of reasons, but it is generally a positive sign, one that gives the impression that the smiler means well, and might be inclined to cooperate (Scharlemann, et al., 2001). The Purple agent has a complement of markers that, like smiles, are dependent upon the possessor’s mood. More precisely, these markers indicate that the agent is currently Desirous.

This seems like a better sort of marker, one that leads the agent to, not only an appropriate mate, but to one who is moreover ready to mate. However, agents in these simulations do not require that the mate be Desirous in order to reproduce. It is not that a ready mate is in any way better for mating. Even without this reliance, the presence of this sort of marker does have significant repercussions on the evolution of coupling strategies. This makes it especially interesting, for it hints that the evolution of displays of sexual readiness may precede differences in receptivity, and be an important aspect of a mating system even without.

Because of the many purposes of smiles, it may not seem an ideal example of a Desire-dependent indicator. Moreover, everyone smiles, so it might not seem an obvious target as a source of sex differences. There are better examples, though, of markers that are both less ambiguous than smiles, and wholly sex-specific. Consider, for example, a pheromone that only shines under certain hormonal conditions. Such a marker might accurately or exclusively indicate a readiness to mate. Perhaps the clearest example of such markers are the lights of fireflies, which they only fire when ready (Lloyd, 1981). There is some evidence that female quail become more attractive to males when sexually receptive, which is suggestive of such a marker in quail hens as well (Bowers, 2000). Such markers may be important cues in many mating systems.

Implementing Smile-markers in Purples

The Desire-dependent markers are like the markers described under Beauty and Ugliness, except that these can only be sensed when the agent is ready to look for a mate – that is, when the agent is currently Desirous. As in previous additions to the model, the Desire-dependent markers come in three sorts: one specific to each sex, and one that appears in both. They occupy markers 4 to 6, with Marker 4 being unisex; Marker 5 exclusive to Blues; and Marker 6 exclusive to Reds. As with other markers, for the sake of most tests they are permitted to be either attractive or aversive (or both) to other agents. The attractions and repulsions to them are identical as to other markers.

Experimenting with Smiles

When the Desire-dependent markers 4 to 6 are run alone, they yield results very similar to those from the unconditional markers of the previous section. Attractions develop toward markers that manifest in the opposite sex; repulsions develop toward markers that manifest in the same sex. This happens in both sexes. The agent in this situation is in a conundrum, for the only way for it to be found by a mate is to be Desirous itself. And of course, it must be Desirous to find a mate. This enforces high Desire in both sexes, and so the symmetrical outcome in the Coupler's Dilemma.

When the three unconditional markers of the previous section are added to this scenario, the results are rather different. Here it is seen that the smile-marker is maladaptive, for all three such markers subside under competition with the unconditional markers. They are selected against until they have no influence on the agents' behaviour. The fact that such apparently maladaptive attractions readily evolved when they were the only markers available suggests that a much stronger quality that contributes to a marker's likelihood to take hold is the degree that it correlates with sex. A stimulus that correlates well with the opposite sex can evolve to become an attractant even though it may also correlate with a maladaptive quality.

Smiles spark the evolution of the asymmetrical outcome of the Coupler's Dilemma in two sorts of simulations. The first is with a lopsided mixture of markers, with the smile-marker in only one sex, and one or more unconditional markers in the other sex, with or without sex-neutral markers. Smiles in this situation encourage the asymmetry by demanding only Desirous callers, and at the same time, foiling the smiler's own mating attempts. Commitment to defection is enforced. The way this happens is explained in the following paragraphs. The other case is when the three Desire-dependent markers are added to the already skewed situation in which other markers in one sex are more frequent than those in the other.

What is happening?

Everyone being able to tell whether you are in the mood or not has some queer effects on coupling. Just when you might have acted, you become especially attractive. And the same goes for everyone else. The overall effect is a bias for mutual cooperation, with cooperators matching with cooperators. With cooperators occupying each other, less defectors get to mate. It seems to be heading toward a stable outcome of symmetrical investment in coupling. Yet, it has an unbalancing effect.

By possessing such a marker, the sex is placed under a selection pressure for higher Desire scores. This is because agents are more attractive when Desirous, when the conditional marker shines, and so those that are Desirous more often, will be attractive more often. As it is now a matter of beauty, its selection is to the other sex's benefit, not the possessor's. Members of the sex will be selected for higher sexual energy, just because it makes them look better.

What of the other sex? If they begin to do the same, they get stuck. The problem is that the Desirous are now both the motivators of sex, and the preferred objects of sex. The desirous other sex is now motivating most of the sex, and some part of them are preferring those that are more attractive on account of their intention to mate in the current round. But an agent can only reproduce once each round, and just when they are in the mood to do so, they are foiled by being the object of another agent's advances. The agent's success that round in reproducing, then, had nothing to do with what it was attracted to. What is the value of being Desirous if it does not lead the agent to initiate

mating? Recall that Desire is a hindrance to survival. If agents are to front the cost of mating they should prefer it not to have been redundant. Further, the very agents the Desire-dependant attraction leads one to try to mate with are the ones most occupied with mating with others, and so many an attempt will go in vain.

This conundrum applies to both sexes, but when Desire levels in one sex begin to be greater than in the other, the desirous sex begins to escape the cycle. As the less desirous sex marshals a smaller army of mate-seekers each round, it tends to be wholly occupied with mate-seekers of the other sex. This reinforces the rise in Desire in the desirous sex, for there is little chance for such an agent to reproduce except by seeking opportunities. Desire scores soar as non-seekers leave no fry.

Meanwhile, in the less desirous sex, displaying one's desire is losing its appeal. This is because an agent of the desirous sex, that prefers to mate with a less popular partner without the marker, begins to have a much greater chance of reproducing. Further, their own efforts to mate are more often leading to redundancy, while the less desirous are harvesting more, and so surviving longer. And most importantly, the other sex is heading in the other direction, towards higher Desire. As average Desire scores in the other sex rise, more are seeking each round than there are Desire-displaying partners to occupy them. This leads them to a selection pressure to drop attraction to the Desire-dependent marker, to focus on the now more available un-Desirous individuals, who will therefore begin to succeed more often in mating.

Now that one sex has committed to cooperation, given the stable asymmetrical solution to the Coupler's Dilemma, the other has a ticket to defect. And it does. This sort of attractant, in this way, carries the consequence of unbalancing the evolution of coupling into sex-specificity.

That the smile-marker, as described here, appears to be maladaptive in the Purple's situation suggests that such releasers might be so elsewhere. Of course, this will not lead us to deny the existence of smiles but it may help to clarify what they mean. They apparently are not indicators of sexual readiness, as that applies to the Purple agent. Smiles may yet be a part of sexual signalling—they may, for instance, indicate sexual receptivity—or be only indirectly related to sex, as a general indicator of cooperation. In either case, it is instructive to note that the Purple agents have no good reason to smile.

Cool Looks

When people pose for a camera, they might smile, but when they are especially keen to look their nastiest, they do something else: the cool look. This is the look that fashion models wear perpetually, as well as Marlon Brando, James Dean, Superman, Elvis, Nina – the very coolest among us. Does it work? More specifically, what does it do to the dynamics of coupling in Purples?

Just as a marker can indicate that its wearer is in the mood, another sort of Desire-dependent marker is possible, one that displays when the agent is set on defecting on coupling, in favour of feeding. Apart from the human cool look,

there is little known precedence for active cool-look-markers contributing to mating in natural organisms. However, this does not mean that they do not exist, for it would seem a difficult thing to identify in another species: the *lack* of interest in mating. This is not the same as *not* displaying the smile-marker, for selection does not work on the absence of genes (!). Both the hint that the Coupler's Dilemma would encourage conditional attractions, and the results of including them in the simulations, suggest that cool-look-markers may appear more often than they reveal themselves to scientists of other species.

So what is happening in the human cool look? Most people attribute pretence to the conspicuously cool – it is not that they are uninterested, but that they are *pretending* to be. Why pretend to be uninterested when one is not? It must be attractive to be uninterested. And it seems that it is possible to fake this look. When enough people fake a given countenance, when trying to be attractive, and if this is, moreover, often a success, we should expect that we may be toying with a releasing stimulus. There are some reasons for expecting the evolution of such markers which are actually quite close to why it happens in Purples.

Implementing Cool-look-markers in Purples

Whether it is pretence or not in humans, a cool look on a Purple indicates a sincere lack of Desire. The cool-look-markers in the Purple agent are, like the smile-markers, dependant on Desire. The possessor of a cool-look-marker will

display it only on rounds when it is un-Desirous, when it is choosing foraging over coupling. Again, there are up to three such markers, depending on the specific test, corresponding to the three that allow sex-specificity to be manipulated. These are otherwise like all other markers, as are the attractions and repulsions to them.

Experimenting with Cool looks

When markers 4 to 6 are converted to Cool-look-markers, and run alone, they behave much like the unlinked markers, Marks U, R, and B. Opposite-sex cool-look-markers evolve to become attractants; same-sex cool-look-markers evolve to become repellents. The sexes share the costs of coupling, and a symmetrical solution to the Coupler's Dilemma is reached. When added to Marks U, R, and B, cool-look-markers again behave no differently. It turns out that whatever combination of these 6 markers are run together—apart from retreating to the artificially unbalanced situation that brings out the asymmetry in even the unlinked markers alone—no asymmetry emerges.

Attraction to cool looks encourages coolness in the other sex. It does so by imposing a selection pressure for low Desire. The attractant, after all, will only make its possessors more attractive when it shines, and that only happens when the agent is ignoring sex. This, in turn, in the spirit of Fisherian sexual selection (see Section 3.8), will leave offspring with the marker, with all its encouragement to be cool. Further, consider that the other parent, the agent who was attracted

by the marker, will always have been Desirous—one of the two must be, and the marked agent apparently was not, at least if it was displaying its cool-look-marker. So the marker tends to lead to offspring whose one parent was a hot Red, and the other, a cool Blue. This sounds like the start of the asymmetrical outcome of the Coupler's Dilemma.

However, simulations reveal that cool-look-markers do not lead to asymmetry. The reason they do not is that they allow agents, when in the mood, to compete effectively as motivators of sex. They provide, in the other sex, the basis for a preference for defectors, and agents find this to their individual advantage. Remember that the agent has already decided to volunteer the costs of coupling when the choice between defector and cooperator appears. It is not to the agent's benefit that the other pay as well. Defectors are simply more available, and so by preferring them, the agent decreases the chances of each mating attempt going in vain. Whereas the smile-markers foiled agents' efforts to mate by especially attracting them to the least available agents, attractions to cool-look-markers direct agents to a segment of the population that tends to be more often available.

So even though cool-look-markers make their possessors more successful at attracting mates when foraging, they also make the other sex's mating efforts more efficient, which in turn makes it easier for both sexes to find mates when they go looking. This increase in efficiency comes about by exploiting the same sort of condition that sex-specificity does, except that here the division is not on lines of sex.

As noted above, it is difficult to see cool-look-markers in other species. It is even hard to imagine experiments that could isolate a marker that indicates an absence of sexual intent. We would doubt it even in our own species if not for Grace Kelly and Kate Moss. Absence of evidence may often be good evidence of absence, but not when there is other reason to think that evidence would be hard come by, as in the present case. These results suggest that attractants conditional upon sexual disinterest may be more common than is apparent.

Smiles versus Cool looks

When placed in the same simulations, the Desire-dependent markers interact with each other. Simulations in which one sex has a smile-marker and the other has a cool-look-marker lead to the smilers motivating all the coupling. This holds under a variety of combinations with other markers, from when it is just these two markers, to when all six markers are active. The asymmetry that emerges is especially stark when the smile is the sex's only marker, for the other sex then has no recourse but to rely on this as the only cue to identify appropriate mates.

One situation in which this does not hold is where the sexes are balanced in terms of Desire-dependent markers. If both sexes are outfitted with a marker of each of the three sorts—smiles, cool looks, and independent markers—the smiles recede. Recall that smile-markers lead to fewer successful matings, so they are left behind. Meanwhile, the cool-look-markers nearly keep pace with the

independent markers, and the situation becomes as if it were just these, which still leads to symmetrical coupling. This result accentuates the observation that cool looks, and attractions to them, are adaptive, and that neither part of smiles are.

Kamikaze Love

Under especially harsh living conditions, the demands of survival may become so heavy that only the most resolute foragers will survive, and even the smallest investment in sex is too much. Agents cannot do both: if they invest in sex, they will starve; if they refrain, they fail to reproduce. The species seems doomed to extinction, for it cannot continue unless individuals both survive and reproduce.

There is yet one solution to this predicament: one sex survives; the other carries the whole load of reproducing. The Purple agent is sexually precocious, meaning that individuals are born mature, capable of mating. Sexual precociousness is not so uncommon in natural organisms. When conditions allow them to survive awhile without eating, come what may, some do. They reproduce as much as they can for the length of their lives without ever bothering to survive. Survival, remember, is hard.

Such conditions bring out an extreme solution to the Coupler's Dilemma, in which one sex is never distracted from foraging, and the other is never distracted from sex, unto starvation. Every influence of the asymmetry is here

starkened, for sexual selection is bare, the only driving force. Beauty in the other sex can reach exaggerated heights quickly.

As far-fetched as this scenario sounds, this sort of asymmetry is not without precedent in nature. In some species, one sex is specialised for mating. Winged male ants, for instance, appear in time for sex, and survive little longer. Though it is not precisely that they do not eat throughout their lives, they are not foragers. Nor are ant communities known for poverty, though. The case of ants suggests that the essential ingredient leading to sex-specialists is that one sex be exempt from the pains of surviving, wherever that exemption comes from. In this case, the one sex receives its exemption from the toils of the other sex.

The Independently Wealthy

The fatal situation of the kamikaze lover is tragic. However, what if, when he fell, he were caught by a grand safety net? Imagine a suitor who, like the kamikaze, had no concerns for life when he came calling, but because something would allow him life for free. He would be in the same position as the kamikaze, but without the tragic end. One example of such an animal comes to mind: the upper class human, the don, sir, pasha, or brahmin.

I am not certain whether I am talking about a time since the emergence of

wealth, or whether easy life could come from being the recipient of very much parental care. If the latter, selection would tend to favour great efforts to care for young unto sexual maturity.

One qualm I have about this scenario is that it should be one-sided. In Sugarland, whatever the abandon of the kamikaze Reds, the Blues have still to struggle to survive. The asymmetry, remember, only emerges when living conditions are harsh. For the analogy to hold, the wealth we imagine rescuing the amorous don must not be available to the other sex. Only if there were a social system to enforce financial dependence in one sex might the kamikaze scenario apply to humankind. Is there such a social system?

Sexual Selection in Sugarland

Sexual Selection

Sexual selection is the sort of selection that comes out of the spread of traits that become associated with sexual attraction. Any feature that makes one attractive will spread as possessors leave broader progenies. Such traits may contribute to their possessors' fitness, entirely because of this association, by making them more attractive and so increasing their ability to mate.

Sexual selection comes to us straight from Darwin, but has since then been fraught with doubt (Roughgarden, Oishi, & Akcay, 2006). Why? Anything that contributes to an animal's attractiveness is expected to be selected for,

whatever it is. This much is uncontroversial. The main qualm about this is why the *other* sex should find it attractive (Ridley, 1995). It should indicate health, or good genes, say critics, but sexual selection is invoked to explain the evolution of characteristics with no such relation to fitness, even those that hinder, such as Darwin's own example, the peacock's tail. It took evolutionary biology some decades to get used to the idea that evolution is not always especially efficient or optimal, so people would not be satisfied by suggestions that the signs that attracted animals to each other had no other purpose, and were loath to accept the same of something so apparently cumbersome as the peacock's tail. Such is the claim of sexual selection.

One explanation for why the attracted might want to be so is known as Fisherian sexual selection (named after its founder, R.A. Fisher, 1930; cited in Dawkins, 1976). The benefit to the attracted is that it will tend to mate with possessors of the attractant, and so leave offspring that also possess it. It will further have done this gene service by leaving offspring that are similarly attracted to the marker, for they will tend to further promote the attractant in the same way. The more ingrained the attractant, the greater the adaptiveness of the attraction to it. The marker and its appeal feed off of each other, each providing the other's impetus. It is circular. It is a positive feedback loop.

Sexual selection is often cast in terms of conflict between the sexes (e.g. Roughgarden, Oishi, & Akcay, 2006), but in a way, opposite-sex agents seem more like allies in this. The real competitors are the same-sex agents. It is within-sex races to be more attractive than one's neighbour that power sexual selection, and that leads to excesses like wattles and racks.

Good Genes

The alternative explanation for attraction is that the attractant indicates higher fitness in its possessors, that it somehow correlates with talent in survival. If an animal with serious life concerns is expected to pursue some marker, it had better signify that the possessor is likely to leave strong offspring. An active marker should then be a demonstration of health. Attraction in health is a genome's trick for seeing beyond itself. A healthy agent is one that, whatever its genes may be, is doing alright in the world, and is likely to leave a similar package to its progeny. If health were attractive, it would allow a selective mater to take advantage of this, and pursue any combination of genes that does well in the relevant world, however the parts would do on their own. So being especially attracted to signs of health in opposite-sex conspecifics is to the selective advantage of an agent's line (Thornhill & Gangestad, 1999).

What sort of marker, though, could accurately show health? "Attracted to health" is not enough, accepting that a gene is unlikely to be able to conceptualise health. So how might a gene simulate a concept of health? Any simple indicator of health could be easily copied, and as soon, throw off the indicator—for instance, fat could mimic muscular; skinny could mimic lithe. Where some agents are better survivors than others, an interesting effect is expected to emerge in the evolution of attractiveness: the Handicap Principle (Zahavi, 1975).

In essence, the Handicap Principle is that indicators of something that impedes survival will evolve to be attractive, for it requires that the encumbered

individual be all the fitter in order to mate *despite* the impediment. Consider, for example, the peacock's tail. It is too long. It impedes flight. It is heavy to carry. If a peacock falls ill, the tail is dragged, and a dirty tale holds no charm. In this way, the handicap is said to be an "honest" sign of health, for it is not something that can be faked. Think of long fur on a cat, how different it looks after a week of easy living, as opposed to during the mid-winter blues in the dirty streets. Is well-groomed, long fur attractive to other cats? Whether it is or not, it is an apt example of the sort of attraction expected by the Handicap Principle. Other examples include indicators of high testosterone (the sex hormone) in development, such as pronounced jaw lines and brow ridges in human males. This is because testosterone both handicaps immune function, and affects facial development over quite a long period of growth (Thornhill & Gangestad, 1999).

The Good Genes advocate would be being unfairly demanding of Darwinian sexual selection if he were to reject it on grounds that it provides no impetus for the attraction. Sure, the marker, the object of beauty, is nothing without the attention of the subject. True too, Darwin provides no explanation for the coming about of this subjecthood. However, the same can be said of good genes. An indicator of good genes, just as an indicator of otherwise empty sex appeal, cannot have its hypothesised effect until there is an attraction to it. Both theories explain the spread of such an attraction, but not its initial appearance.

These two explanations, Good Genes and sexual selection, are competing theories for what lies behind sexual attraction. Whatever the heat of this disagreement, neither excludes the other, so some combination may appear. What do the Purple agents have to say about this matter?

Sexual Selection In the Purple Agent

The results of this study do suggest a stance on this issue. They suggest a compromise, but with sexual selection in the leading role. The main observation is that once an attraction appears, whether to indicate good genes, health, or whatever, sexual selection will wrest attractiveness away from it, stretching the marker to exaggerated heights, from which it can, at best, testify to the extremes of whatever may have once lain behind it.

First note that Purple agents do not vary in terms of survival ability. Each agent harvests sugar as well as any other. They metabolise at the same rate, move at the same speed, sense with the same extent and precision, and so forth. This part of their codes is neither a genetic matter, nor permitted to vary. This means that whatever evolution may occur in the attractiveness and attractedness of Purples is not to be attributed to good genes. The one exception to this is the Desire.

The effect of Desire leading agents to prefer sex over food is a distinct handicap to survival. The Desire-dependent markers, then, fit the description of the sort of markers the Handicap Principle might be expected to power. However, the Handicap Principle can hardly take hold in the current situation. The missing ingredient here is other variable traits to make some agents better survivors than others – something to shine through the handicap in the especially able. Without this, the handicap is just... a handicap. If intact, it could do no other than reinforce the asymmetry that emerged in the simulations. Completing the Handicap Principle scenario is one area to develop in future

work.

As the markers that indicate Desire are exposing a handicap, those that indicate a *lack* of Desire expose a benefit. The Good Genes thesis predicts that agents should develop attractions to such markers more than to others. Note that this is not an ideal trait to test the Good Genes thesis, for its evolution has other, more compelling interpretations. Nor was it designed for such a test. However, it does give some hints here, and at least has to be acknowledged in the present discussion of Good Genes. Such markers and attractions to them do develop beyond those of the handicap marker of high Desire. However, it does not outpace attractions to markers linked to no survival boons at all.

In the Purple agent, sexual selection is clearly occurring. The R-Marks and B-Marks indicate no health, genetic vitality, nothing except that it is another agent, and whether it is male or female. Yet they are both regularly, positively selected for, keeping pace even with markers linked to low Desire, which offers the attracted agent a real advantage.

The stronger argument for the impotence of Good Genes is that health cannot keep up with beauty. The moment some marker begins to become an object of attraction, sexual selection runs away with it, pushing it to whatever heights it can achieve. The genes behind the markers, after all, do not care whether they are healthy to possess or not. Nor should they be expected to. They will exploit any chance attraction to them, wherever that attraction comes from, and for whatever reason it comes. If an attractant reaches some ceiling, beyond which it does not vary, it can clearly no longer differentiate between gradations of health. Even if it keeps evolving, it will soon outpace any differences in health

it may once have indicated.

To illustrate, consider a marker that is linked to an overall measure of health. Say it begins ranging between one and seven, where seven indicates good health, four, moderate health, and one, rather poor health. Enter an attraction to this marker. After several hundred generations, all agents have scores above six. Without mutation, they will soon all have the same score of seven. Permitting mutation, this trait may continue to be variable, but in a short while, the range will grow to be between, say, 18 and 20. Does a 20 indicate greater health than an 18? Not likely, but even if it does differentiate between the very highest degrees of health, continue the simulation for several thousand more generations, until numbers reach 30, 40, or 100. Eventually, sexual selection will push any attractant either to its ceiling, or to meaningless heights, such that it can no longer indicate good genes.

There is one exception to this: the handicap. The reason the handicap can sometimes survive sexual selection with some potency intact is that it is conditional. Hence, the peacock's tail really does lose its appeal when soiled. Note however that this again does not show gradations of strength or health, but merely its presence or utter absence. Those that display the marker are in prime health; those that cannot are out of the running until they do. This is not to say that all markers claimed to be handicaps are so saved. Consider the human testosterone markers, big brow ridges and chins. Are these necessarily constrained to be honest? If not, sexual selection will distort them into lies, and the range of chin sizes will have lost its hypothesised correlation with health. While attractive faces are *perceived* to be healthier than unattractive faces

(Kalick, et al., 1998), research on human facial attractiveness has mostly been at a loss to find a relation between health and features that indicate high levels of testosterone in development, or between health and attractiveness, in general (Shackelford & Larsen, 1999; Kalick, et al., 1998), despite the will to find one (Thornhill & Gangestad, 1999). The Handicap Principle seems to be able at most to explain displays of very gross differences in well being, between being healthy and not.

Why be attracted to a wholly arbitrary marker? The answer is that it never was arbitrary if an exclusive feature of conspecifics, and certainly not if it exclusively indicates opposite-sex conspecifics. By leading an agent to an appropriate mate, a marker has fulfilled the largest part of its chore. This is enough to provide the impetus for Darwinian sexual selection, even without Fisher's addition. In the current work, indeed, there is no firm evidence for Fisherian sexual selection. Trajectories of genes behind the subjecthood of attraction show them to be very weak effects, as if the attracted were undecided about whether they wanted to be so attracted. This is not what Fisher predicts.

Indeed, the factor that was seen to most strongly determine a marker's potential for evolving to become an active attractant or repellent was the extent that it correlated with sex. Any marker that correlates well with the opposite sex is a good candidate for an attractant; any that correlates with the same sex is a candidate for a repellent. Correlation with sex overshadows any other factor, such that a stimulus can evolve to become an attractant even though it may also correlate with potentially harmful qualities. The strength of this effect is underestimated in the sexual selection literature.

One-Sided Love for Better Evolution

Unilateral attraction is not just an empirical possibility, but a good strategy for getting the most out of sex. The point of this section is to raise an issue that any biologist would skip over, because it answers no 'why' questions. The reason it does not is that its apparent function benefits not individuals, not species, but whole orders, so it is hard to think of it as something that could appear on account of the selfish motives of natural selection. To the present study, however, the motive is not to unlock the mysteries of evolutionary biology, but to glean hints about the dynamics of a species caught under such circumstances. So I now take some time to discuss how the unilateral attraction model may promote evolution, how it may be of use in applications of evolution that rely on a model of reproduction that is sexual.

The point is best explained by first considering the alternative: bilateral, fair mating efforts, in which males seek good females, and females seek good males. This leads to the frustrating result that the top quality females end up finding the top quality males, the second quality males and females match; and likewise with third and subsequent quality competitors, whatever the criterion for mate quality may be. That this is what happens was elegantly demonstrated by a very simple simulation study. Upon the foreheads of each of thirty participants was placed a number, such that they could not see it. They were then let be in a room, and told to pair up with the person with the highest number that would accept them. High numbers ended up matched with high numbers; low with low (Ellis, cited in Ridley, 1995).

Now consider what happens under unilateral attraction, where one sex is attractive, and the other attracted. Imagine what would happen in the Ellis study if there were only numbers on the foreheads of the one sex. All Reds would go after the top Blue, but only one could manage. They might fight over him, try to bribe him, or respond in any of the other ways animals do when they find themselves in this circumstance. Though this method is much crueller than the two-way attraction, it affords two important improvements: it keeps the pressure on, and the gene pool mixed.

As with diploidy and dominance, the first issue is the preservation of variability. If mates line up in terms of prettiness, as in the Ellis study, and pretty males always marry pretty females, and ugly males marry ugly females, like is marrying like. The gene hiding action of diploidy is stifled, as more traits become homozygous (i.e. corresponding alleles are similar). Recessive traits are exposed to selection more. The population is bound to become more homogeneous.

Further, if only the fittest males, and only the fittest females reproduce each generation, great amounts of variability will be lost with the masses that fail to reproduce. Consider, in contrast, if only one sex is selected so, as in the unilateral attraction model. In this case, only one sex is under great threat of not reproducing. If the poorest half of these fail to leave offspring, the loss is tolerable, for their opposite-sexed siblings are yet contributing to the gene pool. The only part of the genome that will be especially battered by this is the lustful sex's sex chromosome. Humans are a polygynous species (i.e. where males tend to each marry more than one female, and so where many males will fail to marry even once), and one in which the male sex chromosome only appears in males

(unlike the female sex chromosome which appears in both sexes). Hence, the male sex chromosome is almost wholly receded, apart from its sex determining functions. A martyr has been taken.

Unilateral attraction keeps the pressure on! The biggest fault with bilateral attraction is that beauty is skin deep. Tooth, claw, and even bribes run deeper. As is argued in Section 3.7, beauty is unlikely to indicate health or good genes. However, it can yet indicate absolute mate quality in the other sex. It does this directly, by focusing competition on the same targets, letting only the strong, rich, toothy, or ambitious reproduce. Under many conditions, this situation will lead to polygamy. This will magnify the effect by forbidding reproduction of a potentially large part of the competing sex. For every man who has five wives, four men have none.

The other sex does not need to go through all this for two reasons. For one thing, it would be redundant. The species already has its sieve in the desirous, competing sex. Two colanders stacked hardly strain better than one. Likewise, given that both male and female are required to be behind every birth, having one partner show his mettle is enough to assure good genes in the offspring.

Secondly, though the beauty may only be skin deep, it assures its possessors good mates. By simply being wanted by all, it is implicitly demanding only the fittest competitors. Hence, the beauty tends to be inherited by offspring of the strongest parents. More to the point, it has been assuring its possessors good mates in this way for generations, such that the marker will have become associated with mate quality. The especially beautiful Blue not only assures beauty to any little Blues it may produce, but good genes to both Red and Blue

offspring as well.

However, this only works if the other sex is freed of beauty. Not only does the other sex not need to endure the pains of attraction and competition, it is actually counter-productive. This is because of the greedy spread of sexual selection forever outpacing whatever good gene it might have been linked to. If both are chasing beauty, doomed to be relative, they will end up chasing ever-shifting cascades of fancy in each other, despite their best interests.

Furthermore, the passive sex is not passive about life. While the lusty sex is competing for mates, the less lusty sex is competing for food. If talent in foraging is not a sex-specific trait, as one should scarcely expect it to be, offspring of both sexes are inheriting this parent's talent in harvesting. In this way, such a division of labour allows selection for the more mundane concerns of living to keep pace with the potentially overpowering selection pressures revolving about sex.

What of no attraction at all? Recall that agent populations without any attraction fail to maintain their numbers, and die out, in the simulation environments used.

Another two observations about the action of sexual selection in Purple are worth noting. One is that it is not so hard on variability, even while polarising reproductive success. The pull of runaway selection for attractive traits overshadows other trends. Whereas evolution of the active markers follows straight, inclined lines, those of all other traits, including attractions to markers, are crooked and confused. By having this strong selection criterion overpowering all others, variability is largely hidden from natural selection, in a way analogous

to the effect of diploidy. Once again, this is to the detriment of evolution if it acts in both sexes—after all, natural selection is what powers evolution to solutions—but rejuvenating if restricted to one.

A related observation is that it holds some promise for speciation. Among the overshadowed effects, rolling in the wakes of the markers, are the attractions to them. As these follow the former, they are doomed to always be relatively much weaker effects. Furthermore, they are working against each other. This is because an agent attracted to two signs is going to be led to attempt to mate with possessors of both. That is, it is led to the intersection of two groups—to possessors of Markers A and B, to the exclusion of possessors of either one—greatly narrowing the pool of possible mates. Competitors, likewise, will more often tend toward the intersection (because agents with both markers will attract all those competitors attracted to both, plus some part of those attracted to only one). Therefore, agents attracted to fewer markers (or more completely overlapping markers), will have less mating attempts frustrated by competition, and so will leave more offspring. Selection will tend to push agents into the arms of a single active marker, and these may easily be different markers.

Remembering the positive feedback spirals of sexual selection, how the objects and subjects of attraction tend to encourage the occurrence of each other in the same lines, this situation is likely to lead to groups of related agents mating exclusively with each. That is, it will lead to the emergence of species.

The unilateral attraction model of coupling provides several advantages over bilateral attraction, many of which are not immediately obvious. It is also simpler, and so may be less expensive in terms of processing. Where a modeller

of sexual agents wants good solutions from evolution, they are recommended to consider endowing their agents with asymmetrical attraction, where one sex chases, and the other gives them something to chase.

CHAPTER FOUR

CONCLUSIONS

Even inconspicuous aspects of sex markedly affect the way a species evolves. Simulations of diploidy, allele dominance, and attraction show them each to make interesting predictions about sexual species. Though these are very different aspects of sex, beyond sharing a common source, and so tending to concur, a further theme binds them, that they allow pockets of variability to swell, despite Natural Selection.

Diploidy

Simulations of diploidy and allele dominance show them to act the way classical population genetics predicts. Diploidy retains genetic variability. It does so by simply reserving an extra seat in each part of the genome for what haploidy dismisses as garbage. Because of this, agents with the haploid genome were quicker to reach more optimal trait values, but at the cost of variability, and

hence the potential to continue evolving. Meanwhile, despite their extra baggage, diploid agents reached comparable average heights, with only slight reluctance to approach optimality, and without surrendering their capacity to yet adapt. Looking at a genome as a data structure, this redundancy of diploidy is likened to a form of implicit memory.

In these simulations, as wherever else these issues appear in the evolutionary computing literature, haploidy is plotted against diploidy, and their relative performance is noted, given the circumstance. Haploidy usually does well, but, for instance, where the criteria for the agents' success can change, it is seen to be brittle than diploidy. However, this is looking at the comparison backwards, according to the motives of this study. The present question is not so much which is the better method for which situation. Rather, the intention here is to see what sort of creature either method leads to. Haploidy leads to an organism content to fill the present niche in an optimal way. Diploidy makes an organism less able to commit, and so keeps it on the run to new niches.

Dominance

The effect of Mendelian dominance is again to preserve variability, but according to a different pattern. Dominance relations make more of a difference where variability is high. Where variability is low, as in the later portions of the simulations of Chapter Two, Mendelian dominance and codominance perform similarly. However, where variability is high, Mendelian dominance is seen to

especially buffer selection, slowing the pruning of imperfect alleles. In contrast, codominance is eager to stoke the impatient fires of Natural Selection. The importance of such leaps of faith in codominance is compounded as genomes become more complex, and genetic drift becomes a major player in evolution. In cutting large swathes of variability so quickly, to select against even a very damaging allele, much potential is lost from elsewhere in the genome – potential that might otherwise have become dissociated from the damaging allele, if only it had been allowed to persist for more generations. This is the protection that Mendelian dominance affords. By slowing Natural Selection, it allows ultimately more optimal outcomes, even while preserving variability.

Though such results would not surprise the geneticist, they are not obvious to the rest of us. Certainly, the idea that evolution is bettered by holding Natural Selection at bay runs against some common intuitions. Such intuitions manifest in much of the evolutionary computing literature, as elsewhere, so the point is worth making anew. Further, by going into the matter, one sees more clearly how it works, and it is the subtleties of the natural world that are the more enjoyable to grasp.

Coupling

The approach taken toward the matter of coupling was more explorative. The founding intuition was that the need to couple imposes behavioural constraints on a species, and further that this problem has some features that

may profitably limit a search for how such constraints might lead evolution. To clarify, the matter is the bringing together of the male and female in a two-sex species. The first item to note about this task is that it is a job for one: where two things need to find each other, they will do no better if both look than if just one does. The next observation is that performance of the task comes to the benefit of two. These together open the possibility for a division of labour between the sexes in coupling efforts. Coupling may evolve to be either one-sided or symmetrical. Thirdly, there are not only benefits, but real costs associated with coupling efforts. Being attracted to the opposite sex is not as easy as it seems. Further, these costs are borne wholly by the party that seeks. So the situation is that two parties share the benefits of coupling, but bear the costs individually.

Who will front the costs involved in motivating attraction? There are two sorts of agent under the same selection pressure to couple. So in the context of sexual motivation in one, the other can ignore sex and spend its energies in the selfish pursuit of foraging. By cheating on the common goal of coupling, either sex compels the other to work even harder, and more reliably, than before to make sure that coupling occurs. Hence, we expect one sex to evolve to defect on coupling, leaving the burden on the other. In turn, the consistent generosity of the one allows the more blatant abuse of it by the other. However, this scenario seems to require some cooperation within one sex, analogous to a workers' strike, and so raises some suspicion. It is argued that attraction is one area in which within-sex competition might be overcome, however unlikely elsewhere. It remains to the simulations, though, to demonstrate whether it indeed does or not, and under which circumstances.

This situation in which partners become opponents provides a qualitative criterion by which to analyse simulations of the target species, the very artificial Purple agent of Chapter Three. Under which circumstances will the sexes specialise in coupling, such that one seeks and the other is sought? The motive in sticking to this extreme outcome is merely in recognition of the inexactness of the method, and qualms about what could be thereby generalised. The Purple agent must be modest about what it purports to tell of the natural world.

What makes the difference between whether agents evolve into a pattern of one-sided attraction, or not? Simulating the matter led to several answers to this question. One condition that reliably makes the difference is harshness of conditions. This is a mixture of many population parameters, primarily population size, and relative richness of nutrients, regrowth rates, and such, but can also be expressed in terms of the agents, via metabolic rates. Though this was not a target variable, over several hundred simulations, it became clear that the asymmetry only emerged when living conditions were harsh, and so the cost of coupling sufficiently high. For this reason, all simulations reported upon here are with agents living on the edge of subsistence.

When this dimension is exaggerated agents may come to a particularly extreme, desperate solution. When agents are living beyond their means, the selection pressure to forage may reach such a height that the species cannot afford any investment in sex at all. Under such conditions the species is doomed to extinction unless one sex becomes exempt of foraging. If agents are born sexually precocious, and with enough stores to survive long enough to mate, they respond by effectively jettisoning one sex, such that it does nothing but mate for

its short life, while the other does nothing but survive. Though this is a mockery of the hypothesised asymmetry, surprisingly, it does seem to have precedents in nature.

The presence of mood-dependent attractants also affects the asymmetry. Displays of sexual intent, attractants that are sensed only when the possessor is sexually motivated, place a selection pressure on higher sexual motivation. Such attractants tend to lead agents to fewer successful matings per attempt, for they lead most agents to try to mate with the same portion of the population on each round. This unbalances the situation, leaving it prone to falling into the criterion asymmetry. It also means that responding to such attractants is maladaptive, even while the attractant itself is adaptive, so attractions to them may peter if not held in place artificially.

The other mood-dependent attractant modelled is the very converse: one that is sensed only when the possessor is *not* sexually motivated. Such attractants have opposite effects. They place a selection pressure on lower sexual motivation, and so have a stabilising effect on either outcome of the Coupler's Dilemma. By leading agents to prefer those displaying sexual disinterest, pursuit of such attractants lead to more frequent successes among coupling attempts. It distributes mating efforts more evenly. This means that both the marker and the attraction to it are adaptive. Though hard to identify in natural organisms, this observation that they are adaptive suggests that such conditional attractants may be not uncommon among sexual species.

Beauty can spark evolution toward asymmetry in coupling behaviour. The evolution of the objecthood of attraction is so much stronger than that of the

subjecthood of attraction, that it overshadows. This makes competition to be more beautiful a viable, cost-free alternative to competing to be more attracted or more motivated, on the condition that there is plenty of attractedness and motivation in the other sex to fuel it. Beauty allows the one-sided solution by providing a strong currency for selection.

Another factor seen to be a necessary condition for the emergence of the asymmetry is high variability. Any configuration of parameters that leads to the asymmetry can be reversed by lowering mutation rates. Considerations of patterns of variability combine with variation in beauty in interesting ways.

Attraction and Variability — Distraction and Drift

As promised at the close of Chapter Two, paired with a generous mutation rate, the diploid-Mendelian method of inheritance keeps the Purple agent amply supplied with genetic variability, as much as natural selection can consume, indeed more, such that differences accumulate in all but the most doggedly acted-upon traits. Further, it turns out that the preservation of variability has deeper effects than just keeping the species able to adapt to changing situations, but can influence interactions between the sexes in more direct ways. This point dawned in the context of variable attractants (Section 3.3). When attractants were rarer in one sex than in the other, variability had divergent effects, leading to qualitative sex differences.

Greater variability promotes sex-specificity in two ways. The analogy of a strong current is apt, for it is the pull of genetic drift that increases with variability. As a heavy current will carry away an undermanned boat, and yet be navigable by a purposive crew, so heavy genetic drift will carry away weaker pressures, while letting stronger ones pass. The weaker a selection pressure, the more this sort of effect will act to foil it. This means that wherever a sex difference begins to emerge, for whatever reason, the current will encourage the divide by differentially disrupting the weaker of the two. Where the trait in question involves a division of labour of some necessary task, such as the chore of coupling, the trait will have no less pressure to evolve, and so with the one sex kept swamped in genetic drift, the trait will increase all the more in the other. Such effects come with increases in variability such as when mutation rates are raised.

Beyond simple hikes in genetic variability, beauty further inclines a species towards sex-specificity by raising genetic drift *separately in each sex*. Runaway sexual selection on traits in the hen keep genetic drift high only in her. Likewise, the wattle only boggles the cock. These forces might counteract each other, and keep currents high in both sexes. But as argued above, these tend to lead to asymmetries in beauty in many circumstances. Even where it applies to both, in an imperfect world there will be plenty of opportunities for differences to emerge, such as when a tail becomes too long, or a call too shrill for safety. Were one runaway attractant to begin to press against a ceiling, there would be a commensurate loss in variability, which might be enough to break a balance. For example, suppose the peacock's tail could grow no longer, or more vibrant, and

all were strong enough to wave it high with full swagger. Whatever sway such displays may have on peahens, without some being more or less, there are no grounds for differentiating between them. Discerning peahens may yet rely on other cues to select between cocks, but in any case, the males are freed from selection on the basis of the tail. The pull of genetic drift upon males is alleviated, without affecting levels in females.

The one-sided solution, in turn, leads to different patterns in the distribution of variability between the sexes. In particular, simulations showed one sex transformed into something of a reservoir for genetic variability, exempt of all selection pressures except for beauty, which—to the agents here at least—is wholly superficial. It is these great stores of genetic variability that allow this transformation, not by opening doors through diverse solutions, but by stifling the encumbered sex from reaching an optimal compromise between actions.

To recapitulate, beauty provides a currency for selection, with the potential to overshadow selection for other qualities. This puts selection on a sex-specific treadmill by occupying it with pressures for otherwise superficial features. Each such treadmill in a gene pool raises the sway of genetic drift in a sex-specific manner. High genetic drift swamps weaker selection pressures wholly, discretizing the matter, turning a slight selection pressure into none at all. In cases where a trait is costly, such as coupling, even though a low value may be optimal, the potential to hold on to that optimal value is lost.

Relative Beauty and Good Genes

Good Genes theories fail to explain an innate, releasing stimulus sort of beauty. Sexual selection will take hold of any indicator of attractiveness, and force it to heights beyond sight of its zero, and with it, any correlation with quality. When this happens, beauty is relative. Relative beauty can at best differentiate presence or absence of health, not gradations of it. This means that beauty should not be expected to be a sign of health in its own right. Good Genes concerns may yet lie behind repulsions, for these are not prone to runaway sexual selection.

Next Week's Simulations

The smile- and cool-look-markers are not the only conditional strategies that might have been tried. Conditional complacency, where success with sex determines one's coolness, rather than vice versa, makes sense. Passivity is appropriate for the pretty, but hazardous for the homely. If an agent finds itself swarming with eager suitors, let it not bother chasing them. If there are no callers at all, the agent had better start looking. One can see how such a strategy would serve the agent well: both its pretty and its ugly offspring would follow the best course for them, given their situations. Entry of such conditional traits might foil the hegemony of the beautiful on reproduction—the ugly would have recourse—and so have some interesting effects on the dynamics of a sexual

system.

What happens when beauty peaks? Section 3.3 deals with attractants that never run out of variability. However, peacock tails are not miles long; wattles and jaw lines do not drag on the ground; birdsong is not deafening; and musks are tolerable. Beauty has limits. Further, limits are easy to simulate.

Why Sex?

The question, “why sex?”, has a new answer. How can sex compete with asex, the latter being so much more efficient? This has long been a puzzle of evolutionary biology (Ridley, 1995). One answer is that sex is good at equipping organisms for new niches. So it is not that sex is out-competing asex, but actually the converse, that it only does well when *escaping* from competition with asexual organisms. Asexual things hold their own niches, forcing sexual ones into diversity, which sex makes them a little better at. They are good at diversity because of diploidy, allele dominance, and the mix that comes with coupling. (The point is summarised in the Appendix).

This is a nice explanation for how sexual organisms strayed so broadly, and became so complex, but the question remains, for what is now keeping these complex organisms from lapsing back into asex? The first rejoinder to come to mind is that they have simply inherited sexual ways, and old habits are hard to break. This fails by itself, though, for there are many species that are sexual facultatively, meaning that they can switch between sex and asex. My speculative

answer to this, from considering the results of this work, is that there is a second surge in diversity among organisms that made a firmer commitment to two distinct sexes, and that this is a harder habit to quit. Aphids and rotifers must have their own reasons for indulging in occasional sex. There is no dearth of explanations for why they might. However, no mammal has recourse to an asexual cycle; nor do other of the the most complex organisms, perhaps because they have inherited, in addition to sexual ways, the ways of strict maleness and femaleness. We are sexual because it took sex to accumulate so much diversity; we are stuck being sexual because we are products of a commitment to two sexes, from which there was no turning back. Evolution can reach farther with a diploid genome than it could under haploidy; it can reach farther still when there is strict dissociation between the sexes, enforced by sexual attraction, in the way demonstrated here.

Conclusion

Diploidy, dominance, and attraction combine to manipulate genetic variability in interesting, subtle ways. Throughout, the strategy of smuggling genes past the impatient pruning of natural selection appears as a recurrent theme. Diploidy does this by retaining a garbage collection; Mendelian dominance, by keeping half of the genes strictly silent; attraction, by flooding the gene pool with strong, arbitrary selection pressures, thereby keeping the species wallowing in genetic drift. And this is all the stronger where the sexes diverge. In was not by oversight that the Serpent slipped the apple only to Eve.

GLOSSARY

asex: a system consisting of the sum of all aspects of organisms that subserves reproduction by a single parent (as cell division, spore formation, fission, or budding, in Carbon-based organisms). Contrast with sex, bearing in mind that it is not just an absence of sex, as the word implies (making it etymologically queer), but a distinct system.

coupling: The finding and approaching of an appropriate mate with such designs.

object of attraction: the stimulus that elicits attraction.

objecthood of attraction: the quality of being an object of attraction; attractiveness.

subject of attraction: the reaction to the eliciting stimulus.

subjecthood of attraction: the quality of being affected by an eliciting stimulus; attractedness.

APPENDIX: The Argument for Sex from its Preservation of Diversity

The fable of “The Tortoise and the Hare” is familiar to most, but in a distorted form, not accurate on a number of details. To begin, there was no race, as such. Though the two protagonists began at the same place, their paths were independent of each other. Further, it was not a hare, but a rabbit, and where he was in a rush to be was on holiday. He had heard of the Bahamas, and would not be kept from it. Meanwhile, the turtle (never a tortoise) had heard the same news, and had a similar plan – which is apparently what spawned the myth of it being a race. The rabbit ran to the Bahamas, as planned, and raised a large family. The turtle was somewhere by Raleigh when he heard of the charms of the Adriatic, and he changed course. Halfway to Dubrovnik he changed his mind again in favour of Thailand, which he had not actually reached when he decided instead to go to Sri Lanka, and is now somewhere in the middle of North Africa on his way to Cape Town.

Who won the race? The rabbit’s progeny has now wholly dominated a small, comfortable Caribbean island. The turtle is sweating the Saharan sun, chapped, spent, forever on the edge of death. He left far fewer fry along his way than the rabbit in all his comfort could. But he left them in Italy, and Tibet, and Ethiopia, and New England – it is they that now populate the world. Who won the race? If it was a race, indeed, the rabbit won. But I told you, it wasn’t a race!

REFERENCES

A-Scape. Brookings Institution, <http://www.brook.edu/es/dynamics/models/ascape/>

Axelrod, R. (1984). *The Evolution of Cooperation*. New York: Basic Books.

Bowers, R.I. (2000). Erkek bıldırcınların cinsel tercihleri üzerinde öncül bir etki olarak cinsel sergileme. *Katarsis*, 3, 7-14.

Branke, J. (1999). Memory-enhanced evolutionary algorithms for changing optimization problems. In Congress on Evolutionary Computation (CEC'99), IEEE, Band 3, 1875-1882.

Branke, J. (2001). Evolutionary approaches to dynamic optimization problems: Updated survey. GECCO Workshop on Evolutionary Algorithms for Dynamic Optimization Problems, 27-30.

Buzing, P.C. (2003). VUScape: Communication and cooperation in evolving artificial societies. Masters' thesis, Department of Computer Science, Amsterdam: Vrije Universiteit.

Buzing P.C., Eiben A.E. & Schut M.C. (2005) Emerging communication and cooperation in evolving agent societies. *Journal of Artificial Societies and Social Simulation*, 8.

Campbell, N.A., Reece, J.B., & Mitchell, L.G. (1999). *Biology, 5th Edition*. Amsterdam: Addison Wesley.

Cunningham, M.R. (1986). Measuring the physical in physical attractiveness: Quasi-experiments on the sociobiology of female facial beauty. *Journal of Personality and Social Psychology*, 50, 925-935.

Dasgupta, D., & McGregor, D. R. (1992). sGA: Structured Genetic Algorithm. University of Strathclyde, Technical Report no. IKBS-8-92.

Dawkins, R. (1976). *The Selfish Gene*. Oxford: Oxford U. Press.

Eiben, A.E., Nitschke, G. & Schut, M.C. (2005). Comparison of reproduction schemes in an artificial society for cooperative gathering. AISB Socially Inspired Computing: Engineering with Social Metaphors Symposium (AISB-SIC 2005).

Epstein, J. M., & Axtell, R.L. (1996). *Growing Artificial Societies: Social Science From the Bottom Up*. MIT Press.

Fogel, D.B. (2000). *Evolutionary Computation: Toward a New Philosophy of Machine Intelligence, 2nd Ed*. New York: IEEE Press.

Gao, Y. (1998). An Upper Bound on the Convergence Rates of Canonical Genetic Algorithms. *Complexity International*, 5.

- Ghosh, A., & Dehuri, S. (2004). Evolutionary algorithms for multi-criterion optimization: A survey. *International Journal of Computing and Information Sciences, 2*.
- Goldberg, D. E., & Smith, R. E. (1987). Nonstationary function optimization using genetic algorithms with diploidy and dominance. In J.J Grefenstette, editor, *Proceedings of the Second International Conference on Genetic Algorithms*, 59-68. Lawrence Erlbaum Associates.
- Herrera, F., & Lozano, M. (1996). Adaptation of genetic algorithm parameters based on fuzzy logic controllers. In F. Herrera and J.L. Verdegay, editors, *Genetic Algorithms and Soft Computing*, 95-125. Physica-Verlag, Heidelberg.
- Jaffe, K. (1996). The dynamics of the evolution of sex: Why the sexes are, in fact, always two? *Interciencia 21*, 259-267 <http://www.interciencia.org.ve>.
- JAWAS: Java Artificial Worlds & Agent Societies. Vrije Universiteit, <http://www.cs.vu.nl/ci/eci/jawas/>
- Kalick, S.M., Zebrowitz, L.A., Langlois, J.H., & Johnson, R.M. (1998). Does human facial attractiveness honestly advertise health? *Psychological Science, 9*, 8-13.

Köksal, F., Domjan, M., Kurt, A., Sertel, Ö., Örüing, S., Bowers, R., & Kumru, G.

(2004). An animal model of fetishism. *Behaviour Research and Therapy* 42, 1421-1434.

Kurup, M.M. A study of dominance and diploidy as long term memory in genetic algorithms. <http://www.kurups.org/papers/GA.ps>

Langdon, W.B., Poli, R. (2002). *Foundations of Genetic Programming*, Springer.

Langlois, J.H., Kalakanis, L., Rubenstein, A.J., Larson, A., Hallam, M., & Smoot, M. (2000). Maxims or myths of beauty? A meta-analytic and theoretical review. *Psychological Bulletin*, 126, 390-423.

Langlois, J.H., & Roggman, L.A. (1990). Attractive faces are only average. *Psychological Science*, 1, 115-121.

Leung, K-S., Duan, Q-H., Xu, Z-B., & Wong, C. K. (2001). A new model of simulated evolutionary computation-convergence analysis and specifications. *IEEE Transactions on Evolutionary Computation*, 5.

Lewis, J., Hart, E., & Ritchie, G. (1998). A comparison of dominance mechanisms and simple mutation on non-stationary problems. In *Parallel Problem Solving from Nature (PPSN V)*, pp. 139-148.

- Lin, S-C., Punch, W.F., and Goodman, E.D. (1994). Coarse-grain parallel genetic algorithms: Categorization and new approach. In 6th IEEE Symposium on Parallel and Distributed Processing. IEEE Computer Society Press.
- Lloyd, J.E. (1981). Firefly mate-rivals mimic their predators and vice versa. *Nature*, 290, 498-500.
- Mauldin, M.L. (1984). Maintaining diversity in genetic search. Proceedings of the National Conference on Artificial Intelligence (AAAI-84), 247-250.
- Merriam-Webster (2000). *Merriam-Webster's Unabridged Dictionary*, Version 2.5. Merriam-Webster Incorporated.
- Ng, K. P. & Wong, K. C. (1995). A new diploid scheme and dominance change mechanism for non-stationary function optimization. In L. J. Eshelman (ed.), *Proceedings of the Sixth International Conference on Genetic Algorithms*, pp. 159-166, San Francisco: Morgan Kaufmann.
- O'Doherty, J. Winston, H. Critchley, D. Perrett, D.M. Burt, R.J. Dolan. (2003). Beauty in a smile: the role of medial orbitofrontal cortex in facial attractiveness. *Neuropsychologia*, 41, 147-155.
- Osmera, P. (2003). Evolvable controllers using parallel evolutionary algorithms. Proceedings of MENDEL '2003, Bmo, Czech Republic pp. 126-132.

Penton-Voak, I.S., Perrett, D.I., Castles, D.L., Kobayashi, T., Burt, D.M., Murray, L.K., & Minamisawa, R. (1999). Menstrual cycle alters face preference. *Nature*, *399*, 741-742.

Perrett, D.I., Lee, K.J., Penton-Voak, I.S., Rowland, D., Yoshikawa, S., Burt, D.M., Henzi, S.P., Castles, D.L., & Akamatsu, S. (1998). Effects of sexual dimorphism on facial attractiveness. *Nature*, *394*, 884-887.

Perrett, D.I., May, K.A., & Yoshikawa, S. (1994). Facial shape and judgements of female attractiveness. *Nature*, *368*, 239-242.

Potts, J.C., Giddens, T.D., Yadav S. (1994). The development and evaluation of an improved genetic algorithm based on migration and artificial selection. *IEEE Transactions On Systems, Man, And Cybernetics*, *24*.

Rhodes, G., & Tremewan, T. (1996). Averageness, exaggeration, and facial attractiveness. *Psychological Science*, *7*, 105-110.

Ridley, M. (1995). *The Red Queen: Sex and the Evolution of Human Nature*. Toronto: Maxwell Macmillan.

Roughgarden, J., Oishi, M., & Akçay, E. (2006). Reproductive social behaviour: Cooperative games to replace Sexual Selection. *Science*, *311* (5763), 965-969.

- Rudolph, G. (1994). Convergence of non-elitist strategies. Proceedings First IEEE Conference on Evolutionary Computation. Piscataway: IEEE, 63-66.
- Russell, S., and Norvig, P. (2003). *Artificial Intelligence: A modern approach, 2nd Ed.* New Jersey: Prentice Hall.
- Ryan, C. (1996). The degree of oneness. In Proceedings of the ECAI Workshop on Genetic Algorithms. Springer-Verlag.
- dos Santos, J.P.P.R. (1996). Universal service: Issues on modelling and computation. DSc. Thesis, Université Catholique de Louvain.
- Schafer, R. (2003). Using a genetic algorithm with diploidy to create and maintain a complex system in dynamic equilibrium. In J.R. Koza (Ed.), Genetic Algorithms and Genetic Programming at Stanford 2003, pp.179-186.
- Scharlemann, J.P.W., Eckel, C.C., Kacelnik, A., & Wilson, R.K. (2001). The value of a smile: Game theory with a human face. *Journal of Economic Psychology*, 22, 617-640.
- Schnier, T., & Gero, J. (1997). Dominant and recessive genes in evolutionary systems applied to spatial reasoning. Australian Joint Conference on Artificial Intelligence, 127-136.

Shackelford, T.K., & Larsen, R.J. (1999). Facial attractiveness and physical health.

Evolution and Human Behavior, 20, 71-76.

Shimodaira, H. (2000). A diversity-control-oriented genetic algorithm (DCGA):

Performance in function optimization. In L.D. Whitley, D.E. Goldberg, E. Cant, L. Spector, I.C. Parmee, & H-G. Beyer (Eds.), *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2000)*, Las Vegas, USA, July 8-12, 2000, p. 366, Morgan Kaufmann.

Simoës, A. & Costa, E. (2002). Using genetic algorithms to deal with dynamic

environments: A comparative study of several approaches based on promoting diversity. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO'02)*, W.B. Langdon, et al. (Eds.), Morgan Kaufmann Publishers, New York, 9-13 July.

Singh, A. (2002). Giving genes their voice: A survey of information expression

mechanisms in genetic algorithms. <http://computing.breinstorm.net/natural+paradigms+gas+competent+genetic/> (last seen on 2006-01-11).

Smith, J.M. (1988). *Games, Sex and Evolution*. Hertfordshire: Harvester Wheatsheaf.

Thain, M., & Hickman, M. (1994). *Dictionary of Biology*, 9th Edition. Middlesex:

Penguin.

Thornhill, R. & Gangestad, S.W., 1999, Facial Attractiveness. *Trends in Cognitive Sciences*, 3, 452-460.

Tinbergen, N. (1965). *Social Behaviour in Animals: With Special Reference to Vertebrates*. New York: Wiley.

Vekaria, K., and Clack, C.: Haploid genetic programming with dominance. Departmental Research Note (RN/97/121) (1997)

Yilmaz, A.S. and Wu, A.S. (2002). A comparison of haploidy and diploidy without dominance on integer representations. In the *Proceedings of the 17th International Symposium on Computer and Information Sciences*, Orlando, FL, October 28-30, 2002, pp. 242-248

Zahavi, A. (1975). Mate selection—a selection for a handicap. *Journal of Theoretical Biology*, 53, 205-214.

FAKE SEX:
DIPLOIDY, DOMINANCE, COUPLING, ATTRACTION

Robert Ian Bowers

Submitted to
The Institute for Graduate Studies in Social Sciences
in partial fulfilment of the requirements
for the degree of
Master of Arts

Boğaziçi University
2006

An abstract of the thesis of Robert Ian Bowers for the degree of Master of Arts
from the Institute for Graduate Studies in Social Sciences, April 2006

Title: Fake Sex: Diploidy, Dominance, Coupling, Attraction

What can be understood of a species when we know of it only that it is sexual? Two seemingly innocuous aspects of sex are considered. In Chapter Two, the matter is diploidy and allele dominance; in Chapter Three, it is the need to couple. The effects of each of these, upon the evolution of a species, is investigated by simulating them in an arbitrary sexual species within the JAWAS agent-oriented simulation framework. The aim is to make predictions for the behaviour of organisms that evolve under these circumstances: with a diploid genome, preserving dominance relations, and under the sway of the social dilemma imposed by the need to couple. Diploidy fulfils its promise of preserving variability, whereas haploidy quickly commits its possessors to the current niche. Allele dominance too preserves variability, and without sacrificing adaptivity. These results echo consistent findings in classical population genetics. The expectation of a sex-based division of labour in coupling efforts leads to the identification of several conditions under which such asymmetries are likely to emerge. In particular, sexual selection upon variable attractants, and the presence of mood-dependent attractants are implicated as ingraining sex-specificity in attraction and sexual motivation in predictable ways. Implementations of these aspects of sex markedly affect how agents evolve. Apart from making predictions about natural sexual systems, these results lead to concerns of efficacy, and so make recommendations for how an artificial system that relies on evolution could benefit from implementing sex.

Sosyal Bilimler Enstitüsü'nde Yüksek Lisans derecesi için Robert Ian Bowers tarafından Nisan 2006'te teslim edilen tezin kısa özeti

Başlık: Fake Sex: Diploidy, Dominance, Coupling, Attraction

Sadece cinsiyet sahibi olduğunu bildiğimiz bir tür hakkında ne anlaşılabilir? Cinsiyetin iki yönü ele alındı. Bu tezin ikinci kısmında diploidy ve alel baskınlığı, üçüncü kısmında ise eşleşme ihtiyacı incelendi. Bu konuların bir türün evrimi üzerindeki etkisi benzetim yapılarak incelendi. Benzetim ortamı olarak etmen-tabanlı bir sistem olan JAWAS kullanıldı. Amaç, baskın ilişkiler içeren diploid gen ve eşleşme ihtiyacı tarafından dayatılan sosyal ikilem koşulları altında gelişen organizmaların davranışları hakkında tahminlerde bulunmaktır. Diploidy çeşitliliği koruma sözünü yerine getirirken, haploidy kendini hemen mevcut hücre sahiplerine bırakıyor. Alel baskınlığı da adaptasyonu feda etmeden çeşitliliği koruyor. Bu sonuçlar klasik nüfus genetiği ile tutarlılık içeriyor. Eşleşmede cinse dayalı bir iş bölümü oluşması beklentisi, asimetric durumların ortaya çıkabileceği bazı durumların tanımlanmasını gerektirmiştir. Özellikle, değişken çekiciler üzerine cinsiyet seçimi ve ruh haline bağlı çekicilerin varlığı ile tahmin edilebilir cins özgüllüğü ima edilmektedir. Cinsiyetin bu yönlerinin gerçekleşmesi etmenlerin gelişimini belirgin şekilde etkilemektedir. Bu sonuçlar, doğal cinsel sistemler hakkında tahminlerde bulunmanın yanısıra, evrime dayalı yapay bir sistemin daha verimli olmak için cinsiyet yapısından nasıl faydalanacağına dair önerilerde bulunmaktadır.

ACKNOWLEDGEMENTS

There were several contributions to this work that I must acknowledge. I wish to thank:

Emre Sevinç for translating my descriptions of diploidy into Java, and installing them into the agent code at a time when it all looked very alien to me. This was no small chore, for it required him to first understand the way the system works. Emre also fixed the data collection capacity of JAWAS, while its developers turned deaf, without which I would not have been able to start.

Özge Büyüksu for translating the abstract into the Turkish.

Levent Akin for sending me a very useful article that I would never have found on my own.

Anonymous reviewers of ECAL 2005 (European Conference on Artificial Life) for very useful comments on an early draft of what matured into Chapter Two, and those of ESAW'05 (Sixth International Workshop on Engineering Societies in the Agents World) for their comments on an intermediate draft of the same.

The Boğaziçi University bursar for partially funding my passage to the ESAW'05 conference.

The thesis committee members for their patience, accessibility, and charity.

Additionally, my advisors for the same, over a much longer period.

Further, beyond his role as advisor, Stephen Voss for proofreading the final drafts with such care.

Sumru Özsoy for rescuing me from my oversights of administrative requirements, and for her enthusiasm in keeping the cognitive science programme active.

Zeze for keeping me distracted throughout, and for not letting me forget important realities like smelly litter boxes.

Lastly, some inspiration and much comedy was got from the countless tales of the counted exploits of field experts, Sinan and Mohammed.

CONTENTS

Chapter	
I. GENERAL INTRODUCTION	1
Sex is a Lamppost	1
Sex is Social	4
About the Method	5
Brief Outline	7
The Simulation Environment	8
Evolution as a Criterion	9
The Agents: SugarAgent and Purple	10
Two Sexes	12
II. DIPLOIDY AND DOMINANCE	14
Existing EC Literature on Diploidy and Dominance	18
Haploidy and Diploidy	23
Allele Dominance	28
The Simulation Environment and Agents	30
Design	31
Results	33
Applying Diploidy and Dominance	42
Conclusions	44
III. THE COSTS OF COUPLING	46
Coupler's Dilemma	47
Method and the Purple Agent	55
Genetic Drift	59
Desire and Other Dependent Measures	60
Beauty and Ugliness	68
Smiles and Cool looks	83
Kamikaze Love	95
Sexual Selection in Sugarland	97
One-Sided Love for Better Evolution	105
IV. CONCLUSIONS	111
Diploidy	111
Dominance	112
Coupling	113
Attraction and Variability – Distraction and Drift	117
Relative Beauty and Good Genes	120
Why Sex?	121
Conclusion	122
GLOSSARY	123
APPENDIX	124
REFERENCES	125