
Optimality modeling and fitness trade-offs: when should plants become pollinator specialists?

Paul A. Aigner, Dept of Biology, Univ. of California, Riverside, CA 92521, USA (paacek@aol.com).

The assumption that flowers readily evolve specializations for pollination by particular animals has been central to a standard view of pollinator-mediated adaptive divergence in angiosperms. Stebbins' Most Effective Pollinator Principle (MEPP) formalized this assumption in proposing that a plant should always evolve specializations to its most effective pollinator. I argue that the MEPP and its corollaries are unsupported by basic models of phenotypic selection which predict that a plant should evolve greater specialization to a particular pollinator when the marginal fitness gain exceeds the marginal fitness loss from becoming less adapted to all other pollinators. Differences in pollinator effectiveness are neither necessary nor sufficient to predict specialization. Differences in effectiveness certainly can foster floral specialization to the most effective pollinator in some cases, but when adaptation to a relatively ineffective pollinator requires little loss in the fitness contribution of a more effective pollinator, plants may exhibit striking specializations for the less effective pollinator. Recognizing that the effectiveness of pollinators is not tightly coupled to their importance in selecting for phenotypic novelty may resolve the mismatch between floral features that appear to represent clear evolutionary responses to specific pollinators and patterns of flower visitation that often seem generalized. To shed light on agents of selection and the adaptive value of floral traits I argue that we must go beyond measures of pollinator effectiveness to investigate pollinator-mediated fitness trade-offs over a range of floral phenotypes.

Darwin (1862) established a widely embraced tradition of explaining the diversity of floral form in terms of coevolutionary relationships between angiosperms and their animal pollinators (Baker 1963, Stebbins 1970, Crepet 1983, Müller 1996, Grimaldi 1999). This tradition is built upon the assumption that natural selection frequently favors specialization by plants for pollination by a particular type of animal (often at the level of order or family) and specialization by pollinators for obtaining floral resources from a particular type of plant (family, genus, or species). In traditional models of plant speciation mediated by pollinators, different types of pollinators, by virtue of their allegiance to particular floral forms, act as both agents of natural

selection, driving floral divergence, and of reproductive isolation, preventing gene flow between incipient species (Grant 1949, Waser and Campbell in press). Despite the importance of specialization in this scenario, little consideration has been given to the conditions under which plants should be expected to evolve specializations for particular types of pollinators (Johnson and Steiner 2000).

In the past 30 years, three ideas about the evolution of specialization to pollinators have received widespread attention: 1) plants should specialize to their most effective pollinators (the most effective pollinator principle); 2) the relative effectiveness of different pollinators indicates the extent to which a plant is adapted to each (the logical converse of 1); and 3) large differences in the effectiveness of different pollinators are a prerequisite for plants to evolve specializations. Acceptance of these ideas has led to a perceived paradox between the degree to which morphologically complex flowers appear specialized to particular pollinators (i.e., past ecological contexts) and the often generalized visitation patterns observed at present (present ecological contexts) (Ollerton 1996). Specifically, we find that flowers are often visited by a taxonomically diverse assemblage of potential pollinators in addition to those for which they seem specialized, that this assemblage is variable in time, and that pollinators for which the plant seems specialized are sometimes less effective than pollinators for which the plant has no apparent specializations.

I argue that the three ideas presented above about the nature of specialization are inconsistent with basic optimality and quantitative genetic models of phenotypic selection. These models predict simply that a plant should evolve increased specialization for a particular pollinator when the marginal gain is greater than the marginal loss associated with reducing the fitness con-

tribution of all other pollinators. Specialization by this process does not require differences in mean or maximum effectiveness of individual pollinators, and when differences in effectiveness do occur, specialization need not be towards the most effective pollinator. The potential for specialization depends only on the asymmetry in fitness trade-offs associated with adaptations to particular pollinators. When these trade-offs are strongly asymmetric, plants may evolve striking phenotypic specializations to relatively ineffective pollinators. Much of the perceived paradox between floral adaptation in past and present ecological contexts disappears when we realize that the effectiveness of individual pollinators need not be tightly coupled with their roles as agents of selection.

MEPP and the current paradigm

Stebbins (1970) formalized the notion that natural selection would frequently favor specialization in his "Most Effective Pollinator Principle" (MEPP). The MEPP asserts that floral phenotype evolves in response to the pollinator taxon that is most frequent and effective in its pollination. Secondary pollen vectors may retard adaptation to the most effective pollinator but will not destroy the genetic integration of a floral type once evolved. These ideas were adopted by various authors (e.g., Feinsinger 1983, Schemske and Horvitz 1984), and Waser et al. (1996) formulated them in a simple mathematical model with two pollinator types. Whereas the model of Waser et al. (1996) suggested support for the MEPP, the underlying assumptions were restrictive: a rare mutation that increased the effectiveness with respect to pollinator visitation rate or pollen transfer per visit (respectively "quantity" and "quality" components of pollination service) for a particular pollinator type would spread if that pollinator was already the most effective pollinator with respect to the additional component, and provided that the mutation also decreased the effectiveness of the inferior pollinator type by an exactly equal amount. The last assumption is critical; it amounts to an asymmetric fitness trade-off because the net gain from evolving greater specialization to the more effective pollinator exceeds the fitness loss from sacrificing adaptation to the less effective pollinator. Waser et al. (1996) did not discuss how changes in this trade-off could alter their qualitative conclusions.

A corollary of the MEPP is that pollinators must vary in their effectiveness for plants to evolve specialization (Schemske and Horvitz 1984). This corollary has helped fuel studies of relative pollinator effectiveness (e.g., Primack and Silander 1975, Motten et al. 1981, Schemske and Horvitz 1984, Herrera 1987, 1989, Horvitz and Schemske 1990, Eckhart 1991, 1992, Fish-

bein and Venable 1996, Gómez and Zamora 1999). For example, in a study of a neotropical herb, Horvitz and Schemske (1990) found that the mutualists (pollinators and antguards) that were of the highest quality were not the most plentiful, and interpreted this as a constraint on the opportunity for the plant to specialize. Similarly, Fishbein and Venable (1996) found that bees were the most effective pollinators of *Asclepias tuberosa* and cautioned against classifying this milkweed as butterfly-adapted, despite its apparent fit to a butterfly pollination syndrome.

An optimality model

Optimality modeling has long been used to predict how foraging animals (including pollinators) should use food resources (Pyke 1984, Ohashi and Yahara 1999) but somewhat surprisingly has not been used in the converse to predict how plants should use pollinators. I propose a univariate model in which pollination service is a function of the expression of a continuously varying (quantitative) floral trait. Furthermore I assume that lifetime fitness is related to pollination service by a continuously increasing function over the range of phenotypes considered (Parker and Maynard Smith 1990, Roff 1994). Hence the pollination (= fitness) function is identical to those used in other frequency-independent evolutionary optimization models (Parker and Maynard Smith 1990). The fitness function is analogous to a selection differential in quantitative genetics under the assumption that the trait under consideration has high heritability.

Consider a plant population expressing a range of phenotypes with respect to some trait, x , such as flower (corolla) length. When faced with two pollinator types, there will be a net fitness function $w(x)$ relating an index of plant fitness (pollen export or receipt) to the floral trait. I assume that $w(x)$ is a composite of the functions $w_1(x)$ and $w_2(x)$ describing individual fitness contributions of each pollinator type with respect to variation in the trait. In the simplest case, the effect of each pollinator is independent of the other pollinator and $w(x) = w_1(x) + w_2(x)$. More realistically the effect of each pollinator depends on the presence of the other (Thomson and Thomson 1992), and $w(x) = w_1(x) + w_2(x) + i_{12}(x)$, where $i_{12}(x)$ represents an interaction between the two pollinator types. Little is known about the nature of these interactions, so I begin by assuming that $i_{12}(x) = 0$.

What can be said then about the conditions under which natural selection will favor specialization to either pollinator type? Fitness is maximized at a phenotype x^* specified by $dw(x^*)/dx = 0$ and $d^2w(x^*)/dx^2 < 0$. This occurs when $w'_1(x^*) + w'_2(x^*) = 0$, i.e., where $w'_1(x^*) = -w'_2(x^*)$. Thus, the floral phe-

notype should evolve towards specialization as long as the marginal fitness gain from becoming more adapted to one pollinator exceeds the marginal loss from becoming less adapted to the other. In summary, specialization is most likely to occur when there is a strong asymmetry in the trade-off involved in becoming more adapted to one pollinator versus the other.

Graphical analysis

Fitness functions might have various shapes, but even Gaussian functions illustrate that the outcome of selection in the two-pollinator model is not always intuitive. Fig. 1A illustrates the scenario as envisioned by Stebbins. Here a plant population is serviced by two pollinators that differ strongly in maximum effectiveness; the population evolves almost entirely in response to selection imposed by its more effective pollinator. The optimal phenotype is virtually indistinguishable ($+0.01\%$) from the optimal phenotype if only pollinator 1 was present, and the population can be considered specialized for pollinator 1.

When fitness functions are Gaussian (or more generally when they are unimodal and continuous), it is easy to show that the phenotype that maximizes net fitness must lie between the two phenotypes that maximize fitness with respect to each pollinator type. Thus, the optimal phenotype in a two-pollinator environment will always represent a “compromise” between optimal adaptation to either pollinator type. As Fig. 1A demonstrates, this compromise may be operationally indistinguishable from complete specialization.

Fig. 1B illustrates that strong differences in effectiveness need not result in specialization. Here, the optimal phenotypes for pollinators 1 and 2 are separated by a broad region where fitness remains essentially constant because the fitness trade-offs for adapting to the two pollinator types are almost symmetrical. Hence the optimal phenotype is a broad peak roughly midway between the two phenotypes that would represent specialization to one or the other pollinator individually. The population is truly generalized to the two available pollinators, even though at equilibrium pollinator 1 is roughly 4 times more effective than pollinator 2.

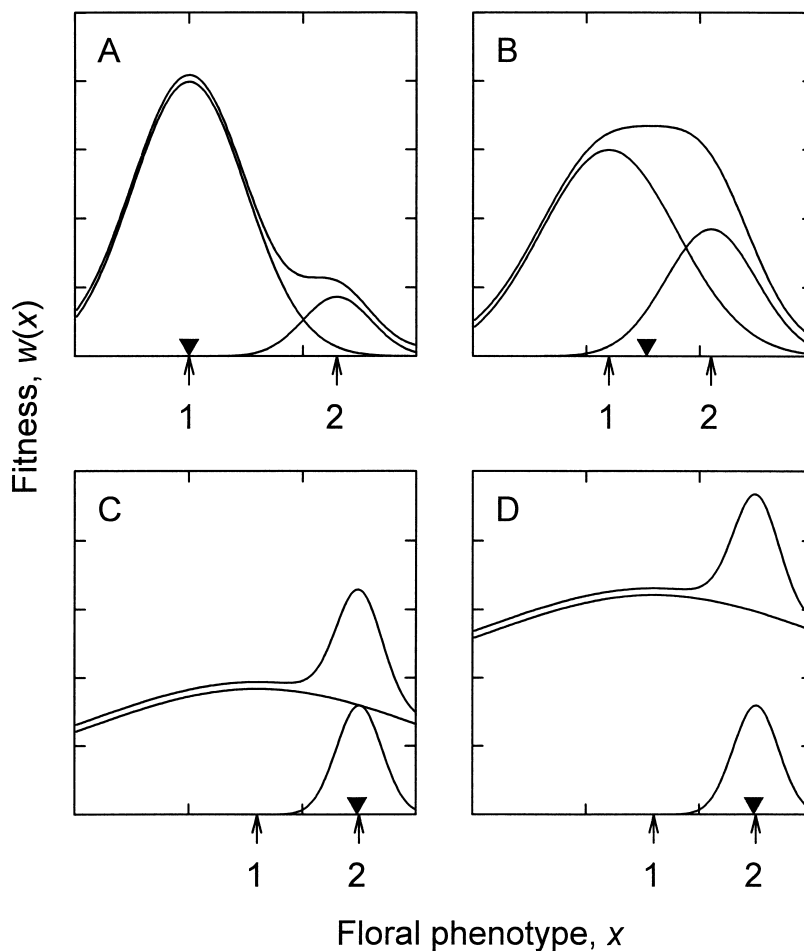


Fig. 1. Four scenarios of pollinator-mediated selection on floral phenotype in a two-pollinator environment, where fitness contributions of pollinators are additive. In each graph the lower two curves represent the individual contributions of two pollinator types, 1 and 2, to plant fitness. The upper curve is the net fitness function (offset slightly upward for clarity) assuming that the effects of the individual pollinators are purely additive. Each numbered arrow below the x -axis indicates the phenotype that is optimally adapted to that pollinator by itself. The solid triangles above the x -axis indicate the optimal phenotype in the two-pollinator environment.

Fig. 1C and D illustrate scenarios that may be the most counterintuitive. Both cases exhibit strong asymmetry in fitness trade-offs for specialization to one versus the other pollinator type. Pollinator 1 is effective over a broad range of phenotypes. Conversely, pollinator 2 is effective on a narrow range of phenotypes. The result of this asymmetry is that pollinator 2 has disproportionate weight in determining the net phenotypic optimum when both pollinators are present. The plant population can afford to evolve specialization to pollinator 2, because this entails little sacrifice in the fitness contribution of pollinator 1. Fig. 1C shows that differences in mean pollinator effectiveness are not required for specialization: at the optimal floral phenotype, to which plants should evolve, the effectiveness of pollinators 1 and 2 is equal, yet the optimal phenotype is virtually indistinguishable from the one that would occur if only pollinator 2 were present. Inspection of Fig. 1C and D reveals that increasing the height of the fitness function associated with pollinator 1 has no effect on the outcome, as long as shapes of the functions do not change. Thus, pollinator 1 could be orders of magnitude more effective than pollinator 2 on all floral phenotypes without altering at all the outcome of specialization to pollinator 2.

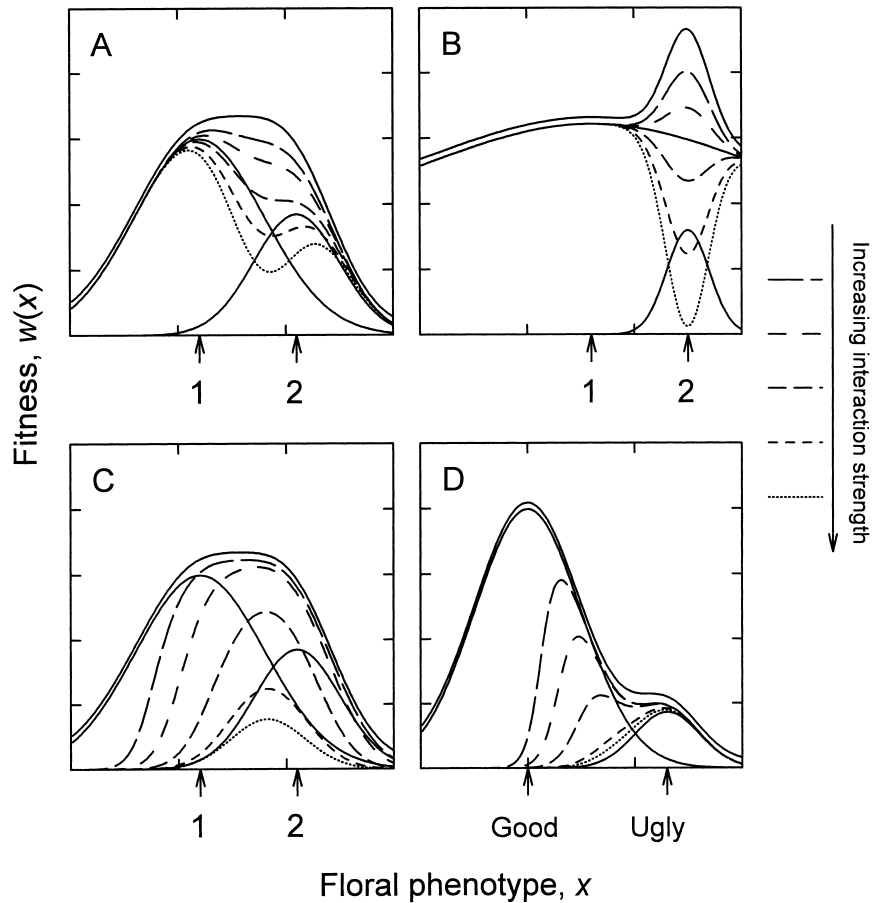
Negative interactions – the good and the ugly

The preceding analysis suggests that, contrary to the prediction of the MEPP, morphological specializations that exclude any effective pollinators should be difficult to evolve; yet nature abounds with examples of floral architectures that guarantee remarkable exclusivity, some of which have clearly evolved from more generalized ancestors. A useful optimality model should therefore suggest mechanisms by which increased specialization can evolve in an environment where several pollinators are effective. Certain types of negative interactions in the fitness contributions of pollinator types may provide such mechanisms. Negative interactions occur when the combined effect of several pollinator types on plant fitness is less than the sum of the contributions that each type would make if it occurred individually. Although there is little empirical evidence for how pollinators might interact in their effects on plant fitness, we might often expect such interactions to negatively affect the quantity of pollen removed and delivered by particular pollinators for the simple reason that pollen is a finite resource; removal of pollen by one species decreases the amount available to be removed by another. The model of Thomson and Thomson (1992) predicts that such interactions are likely to be strongest when “good” pollinators (those that pick up and deposit much pollen) occur with “ugly” ones (those that pick up much pollen, but deposit little of it).

A plausible negative interaction can be included in the two-pollinator model by assuming that, with respect to floral phenotype, each pollinator reduces the fitness contribution of the other by a proportion that is a function of that pollinator’s own effect at that phenotype. In the simplest case, this function is positive and linear; each pollinator has the greatest negative effect on the fitness contribution of the other at the phenotype where its own contribution to fitness is greatest. Such a trade-off increases the likelihood that specialization will be favored in many scenarios. Biologically such an interaction is expected if variation in the effectiveness of pollinators with respect to floral phenotype is due primarily to variation in visitation frequency or to variation in the absolute amount of pollen removed per visit, with the proportion subsequently deposited not changing. This type of interaction can be considered an interaction due to pollinator “goodness” (sensu Thomson and Thomson 1992) because there is no variation in the proportion of collected pollen that is subsequently lost by pollinators. Under this scenario of interaction between two “good” pollinators, negative interactions will be strongest at the floral phenotype where the joint goodness of pollinators is highest, i.e., where both pollinators are visiting frequently or are effectively removing much pollen. The good-pollinator interaction model can be expressed as $w(x) = w_1(x)[1 - \alpha_{12}w_2(x)] + w_2(x)[1 - \alpha_{21}w_1(x)]$, where α_{12} is a coefficient (ranging between 0 and $1/w_2(x^*)$ where $w_2(x^*)$ is the maximum for $w_2(x)$) representing the reduction in the fitness contribution of pollinator 1 attributable to the presence of pollinator 2, and similarly, α_{21} represents the effect of pollinator 1 on pollinator 2. The interaction coefficients are constrained so that the interaction (bracketed) terms represent proportional fitness reductions (i.e., take on values in the interval [0, 1]). The net fitness function is maximized at x^* where $w'_1(x^*) = -w'_2(x^*)[(1 - (\alpha_{12} + \alpha_{21})w_1(x^*)) / (1 - (\alpha_{12} + \alpha_{21})w_2(x^*))]$, demonstrating that the outcome depends only on the sum of the effects α_{12} and α_{21} , and not on their asymmetry.

Whether the outcome of the good-pollinator interaction model differs qualitatively from that of the additive model depends on the overlap between the individual pollinator fitness functions and the strength of the interaction ($\alpha_{12} + \alpha_{21}$). If these fitness functions are partially overlapping, as in Fig. 1A and B, then weak interactions tend to shift the net fitness maximum towards a phenotype that is closer to the phenotype that maximizes the fitness contribution of the pollinator with the highest peak contribution, i.e., increased specialization to the more effective pollinator is favored (Fig. 2A). Increasing the strength of the interaction increases the strength of selection to specialize on the more effective pollinator until, for very strong interactions, the net fitness function becomes bimodal and selection becomes disruptive. If the individual fitness

Fig. 2. Scenarios of pollinator-mediated selection on floral phenotype in a two-pollinator environment, where negative interactions occur in the effects of pollinators on plant fitness. Graphs A and C begin with the additive model of Fig. 1B; graphs B and D begin with the additive models of Fig. 1D and A, respectively. The effects of increasing interaction strength on the net fitness function are indicated by the series of dashed curves. Graphs A and B illustrate the good-pollinator interaction, which is strongest where the pollinator types overlap most in their fitness contributions. Graph C illustrates the ugly-pollinator interaction, which is strongest where the pollinator types overlap least in their fitness contributions. Graph D illustrates the scenario in which a good and an ugly pollinator co-occur. When the good pollinator exerts its maximum negative effect on the ugly pollinator, increasing the strength of the ugly-pollinator interaction shifts the position of the net fitness optimum towards specialization to the ugly pollinator.



functions are completely overlapping, as in Fig. 1C and D, weak interactions do not immediately favor specialization to the more effective pollinator (Fig. 2B); the outcome of specialization to the less effective pollinator is unchanged unless $\alpha_{12} + \alpha_{21}$ is large, in which case specialization to the more effective pollinator is favored or selection becomes disruptive.

A subtle but important effect of strong good-pollinator interactions is that they push apart the positions of the fitness maxima beyond the points where they would coincide with specialization to either pollinator individually (Fig. 2A and B). This suggests that, although weak good-pollinator interactions simply broaden the conditions under which specialization to the most effective pollinator can occur, strong ones may be potent forces for floral divergence, favoring phenotypes even more extreme than those representing complete specialization.

Not all negative interactions promote specialization to the more effective pollinator, nor is the asymmetry in the effects of pollinator types on each others' fitness contributions always inconsequential. "Ugly" pollina-

tors are so named because they deposit only a small proportion of the pollen that they remove (Thomson and Thomson 1992). If the variation in the fitness contribution of an ugly pollinator with respect to plant phenotype is caused by a changing relationship between the amount of pollen picked up per visit and the amount subsequently deposited (i.e., the amount of pollen lost), then the ugly pollinator is unlikely to conform to the assumptions of the good-pollinator interaction model. This ugly pollinator's effect on the fitness contributions of other pollinators will be most negative at that floral phenotype where its own contributions to plant fitness are smallest – precisely at that phenotype where it is wasting the most pollen. To describe the interaction between two pollinators that vary in ugliness, one pollinator's proportional reduction of the second pollinator's fitness contribution should be an inverse function of the first's own contribution to fitness. This ugly-pollinator interaction model can be expressed as $w(x) = w_1(x)[1 - \beta_{12}/(w_2(x) + \beta_{12})] + w_2(x)[1 - \beta_{21}/(w_1(x) + \beta_{21})]$, where β_{12} is a constant (≥ 0) representing the reduction in the fitness contribu-

tion of pollinator 1 attributable to the presence of pollinator 2, and similarly, β_{21} represents the effect of pollinator 1 on pollinator 2. The location of the net fitness maximum now depends on the asymmetry of the interaction. If the interaction is symmetric ($\beta_{12} = \beta_{21}$) and nonzero, then compared to the purely additive model, the net fitness maximum in a two-pollinator environment will always be at a phenotype nearer to the phenotype that maximizes fitness for the less effective pollinator alone (Fig. 2C). The shift becomes greater as the strength of the interaction increases. In contrast to the good-pollinator interaction, the ugly-pollinator interaction restricts the conditions under which specialization to the more effective pollinator will occur.

Pollinator goodness and ugliness are not independent, and interactions in nature are likely to be influenced by both. Because models of the two types of interaction suggest opposing influences in the direction of specialization, it is useful to consider the situation, alluded to earlier, in which a plant is serviced by a good pollinator and an ugly one. The individual fitness contribution of a good pollinator will be represented by a function with a much higher peak than that of an ugly pollinator, because even at its best, an ugly pollinator wastes far more pollen than does a good pollinator. The net form of selection in this good and ugly pollination environment can be specified by the model: $w(x) = w_g(x)[1 - \beta_{gu}/(w_u(x) + \beta_{gu})] + w_u(x)[1 - \alpha_{ug}w_g(x)]$, where the subscripts “g” (good) and “u” (ugly) replace “1” and “2”. Thus the fitness contribution of the good pollinator is reduced by an ugly interaction term and the fitness contribution of the ugly pollinator is reduced by a good interaction term. With α_{ug} at its maximum value, even small values of β_{gu} produce a net fitness function in which the optimum is pulled towards the phenotype that maximizes the fitness contribution of the ugly pollinator alone (Fig. 2D). As tempting as it is to assume that a plant should evolve to a phenotype that excludes an ugly pollinator, the model predicts the opposite – as β_{gu} becomes large, it is easy to evolve almost complete specialization to the ugly pollinator. In effect, the plant must make the best of a bad situation, evolving to a phenotype where the ugly pollinator is least ugly.

The interactions considered thus far represent only a fraction of the possible ways in which one pollinator can affect another’s contribution to plant fitness. Interactions among pollinator types may also be positive if the consumption of floral resources by one pollinator changes the behavior of another in a way (e.g., increased between-plant flight distances, or visits to fewer flowers on the same plant) that increases the quality of pollen transferred. What is clear from the few possibilities considered is that relatively simple interactions may alter fitness trade-offs in complex ways.

Discussion

The apparent paradox between phenotypic specialization and ecological generalization is not unique to flowers or to pollination biology, nor is the solution I have proposed here novel. Cichlid fishes of the African great lakes show a remarkable array of morphological specializations for feeding, yet often prefer readily available food resources for which they are not specialized. Robinson and Wilson (1998) developed an optimal foraging model which showed that morphological specializations for less-preferred food items can allow coexistence when many species in a community share a single preferred food item. They also predicted that such specializations would most likely evolve when morphological specializations required little sacrifice in the ability to obtain preferred food items.

The proposition that angiosperms sometimes evolve phenotypic specializations to relatively ineffective pollinators must of course be tested empirically. One suggestive example is provided by Hurlbert et al. (1996), who found that flexible pedicels in *Impatiens capensis* are an adaptation to hummingbird pollination because they increase flower movement and pollen transfer during hummingbird visits, yet have no effect on the pollination service of bees, which are the primary visitor to this jewelweed. The authors pointed out that the adaptation to an uncommon visitor is possible because there is no cost with respect to bee pollination. Examples of floral phenotypes that seem to represent compromises for two pollinator types are somewhat more common (e.g., Macior 1986, Eisikowitch and Rotem 1987, Sahley 1996, Lange and Scott 1999), but one reason for the current shortage of empirical evidence for scenarios that fall outside the MEPP may be that floral biologists have simply failed to look for it. For example, flowers that possess characters of more than one pollination syndrome are frequently interpreted as being in transition between syndromes (e.g., Macior 1986) or as maintained by temporal variation in selection (e.g., Sahley 1996), rather than as adaptations in their own right. The particular mix of syndrome characters possessed by such a flower may in fact be under strong stabilizing selection because it is optimal under a particular multiple-pollinator environment.

The models considered here suggest that we must move beyond studies that simply rank pollinators by their effectiveness on the mean floral phenotype if we hope to identify agents of selection and floral specialization. A more profitable approach will be to decompose the net fitness function into the contributions of individual pollinators, and to quantify the trade-offs involved in becoming more adapted to a particular pollinator.

Decomposing fitness functions to test hypotheses about floral adaptation in real pollination systems requires experimental manipulations of pollinator assem-

blages. An idealized experiment might consist of a multi-factor design, each factor being the presence or absence of an individual pollinator type, with factors fully crossed to estimate interactions. An ideal study system will therefore consist of a plant with a pollinator assemblage in which it is feasible to individually exclude particular pollinator types. Such manipulations have been performed where pollinators segregate temporally (Morse and Fritz 1983, Jennersten and Morse 1991, Herrera 2000), or where the exclusion technique exploits size differences between pollinator types (Waser 1979, Barthell et al. 1999, Fenster and Dudash 2001), but these studies have been limited in that they measured only mean contribution of pollinator types to plant fitness without regard to floral variation. To quantify fitness trade-offs, the response variable in the experimental manipulation of pollinator assemblages must be the fitness function itself. Thus, experimenters will need to manipulate pollinator assemblages using plants with a range of floral phenotypes. In practice, it will rarely be possible to decompose the net fitness function into the effects of every pollinator species and all of their interactions, but a benefit of an optimality approach is that pollinator types in the model can be defined as broadly as necessary. For convenience, a group of pollinators may be considered a “black box” with respect to its contribution to the net fitness function of a plant, or an investigator may have an a priori reason to assume that several pollinator species are functionally identical. An optimality approach can be employed hierarchically, first to test whether apparent syndrome features are actually adaptations to higher level taxa (such as bees or butterflies) as is often assumed, and second to test whether the fitness functions due to species within syndrome taxa favor the same phenotypic optimum.

When peaks in observed distributions of the frequency of floral phenotypes do not correspond to the outcome predicted by analyzing trade-offs alone, the discrepancy may be caused by interactions in the effects of pollinators. Different types of interactions may affect the likelihood of evolving specialization to the most effective pollinator, so specialization to the most effective pollinator may prove to be a general phenomenon if, for instance, the good-pollinator interaction is common in nature. A good-pollinator interaction would be fostered if pollinator types frequently compete for floral resources and thus have reciprocal effects in reducing the attractiveness of certain floral phenotypes. While such interactions seem likely, empirical evidence is sorely needed.

The models presented here join others that draw into question the generality of the MEPP (Waser et al. 1996, Waser and Campbell in press). Plants may in fact often become specialized to their most effective pollinators, but there are as yet insufficient theoretical or empirical grounds to claim this as an organizing principle of floral ecology and evolution.

Acknowledgements – I thank Nick Waser for discussions that contributed to these ideas and for comments that substantially improved the manuscript. Cathy Koehler, Mary Price and Jeff Ollerton provided helpful comments on earlier drafts.

References


- Baker, H. G. 1963. Evolutionary mechanisms in pollination biology. – *Science* 139: 877–883.
- Barthell, J. F., Thorp, R. W., Wenner, A. M. and Randall, J. M. 1999. Yellow star-thistle, gumplant, and feral honey bees on Santa Cruz Island: a case of invaders assisting invaders. – In: Browne, D. R., Mitchell, K. L. and Chaney, H. W. (eds), *Proceedings of the Fifth California Islands Symposium*. U.S. Minerals Management Service, Pacific OCS Region, Camarillo, CA, pp. 269–273.
- Crepet, W. L. 1983. The role of insect pollination in the evolution of angiosperms. – In: Real, L. (ed.), *Pollination biology*. Academic Press, pp. 29–50.
- Darwin, C. 1862. On the various contrivances by which British and foreign orchids are fertilised by insects. – John Murray.
- Eckhart, V. M. 1991. The effects of floral display on pollinator visitation vary among populations of *Phacelia linearis* (Hydrophyllaceae). – *Evol. Ecol.* 5: 370–384.
- Eckhart, V. M. 1992. Spatio-temporal variation in abundance and variation in foraging behavior of the pollinators of gynodioecious *Phacelia linearis* (Hydrophyllaceae). – *Oikos* 64: 573–586.
- Eisikowitch, D. and Rotem, R. 1987. Flower orientation and color change in *Quisqualis indica* and their possible role in pollinator partitioning. – *Bot. Gaz.* 148: 175–179.
- Feinsinger, P. 1983. Coevolution and pollination. – In: Futuyma, D. J. and Slatkin, M. (eds), *Coevolution*. Sinauer, pp. 282–310.
- Fenster, C. B. and Dudash, M. R. 2001. Spatiotemporal variation in the role of hummingbirds as pollinators of *Silene virginica*. – *Ecology* 82: 844–851.
- Fishbein, M. and Venable, D. L. 1996. Diversity and temporal change in the effective pollinators of *Asclepias tuberosa*. – *Ecology* 77: 1061–1073.
- Gómez, J. M. and Zamora, R. 1999. Generalization vs. specialization in the pollination system of *Hormathophylla spinosa* (Cruciferae). – *Ecology* 80: 796–805.
- Grant, V. 1949. Pollination systems as isolating mechanisms in angiosperms. – *Evolution* 3: 82–97.
- Grimaldi, D. 1999. The co-radiations of pollinating insects and angiosperms in the Cretaceous. – *Ann. Mo. Bot. Gard.* 86: 373–406.
- Herrera, C. M. 1987. Components of pollinator “quality”: comparative analysis of a diverse insect assemblage. – *Oikos* 50: 79–90.
- Herrera, C. M. 1989. Pollinator abundance, morphology, and flower visitation rate – analysis of the quantity component in a plant-pollinator system. – *Oecologia* 80: 241–248.
- Herrera, C. M. 2000. Flower-to-seedling consequences of different pollination regimes in an insect-pollinated shrub. – *Ecology* 81: 15–29.
- Horvitz, C. C. and Schemske, D. W. 1990. Spatiotemporal variation in insect mutualists of a neotropical herb. – *Ecology* 71: 1085–1097.
- Hurlbert, A. H., Hosoi, S. A., Temeles, E. J. and Ewald, P. W. 1996. Mobility of *Impatiens capensis* flowers: effect on pollen deposition and hummingbird foraging. – *Oecologia* 105: 243–246.
- Jennersten, O. and Morse, D. H. 1991. The quality of pollination by diurnal and nocturnal insects visiting common milkweed, *Asclepias syriaca*. – *Am. Midl. Nat.* 125: 18–28.
- Johnson, S. D. and Steiner, K. E. 2000. Generalization versus specialization in plant pollination systems. – *Trends Ecol. Evol.* 15: 140–143.

- Lange, R. S. and Scott, P. E. 1999. Hummingbird and bee pollination of *Penstemon pseudospectabilis*. – J. Torrey Bot. Soc. 126: 99–106.
- Macior, L. W. 1986. Floral resource sharing by bumblebees and hummingbirds in *Pedicularis* (Scrophulariaceae) pollination. – Bull. Torrey Bot. Club 113: 101–109.
- Morse, D. H. and Fritz, R. S. 1983. Contributions of diurnal and nocturnal insects to the pollination of common milkweed (*Asclepias syriaca* L.) in a pollen-limited system. – Oecologia 60: 190–197.
- Motten, A. F., Campbell, D. R., Alexander, D. E. and Miller, H. L. 1981. Pollination effectiveness of specialist and generalist visitors to a North Carolina population of *Claytonia virginica*. – Ecology 62: 1278–1287.
- Müller, A. 1996. Host-plant specialization in western palearctic anthidiine bees (Hymenoptera, Apoidea, Megachilidae). – Ecol. Monogr. 66: 235–257.
- Ohashi, K. and Yahara, T. 1999. How long to stay on, and how often to visit a flowering plant? – a model for foraging strategy when floral displays vary in size. – Oikos 86: 386–392.
- Ollerton, J. 1996. Reconciling ecological processes with phylogenetic patterns: the apparent paradox of plant-pollinator systems. – J. Ecol. 84: 767–769.
- Parker, G. A. and Maynard Smith, J. 1990. Optimality theory in evolutionary biology. – Nature 348: 27–33.
- Primack, R. B. and Silander, J. A., Jr. 1975. Measuring the relative importance of different pollinators to plants. – Nature 255: 143–144.
- Pyke, G. H. 1984. Optimal foraging: a critical review. – Annu. Rev. Ecol. Syst. 15: 523–575.
- Robinson, B. W. and Wilson, D. S. 1998. Optimal foraging, specialization, and a solution to Liem's paradox. – Am. Nat. 151: 223–235.
- Roff, D. A. 1994. Optimality modeling and quantitative genetics: a comparison of the two approaches. – In: Boake, C. R. B. (ed.), Quantitative genetic studies of behavioral evolution. Univ. of Chicago Press, pp. 49–66.
- Sahley, C. T. 1996. Bat and hummingbird pollination of an autotetraploid columnar cactus, *Weberbauerocereus weberbaueri* (Cactaceae). – Am. J. Bot. 83: 1329–1336.
- Schemske, D. W. and Horvitz, C. C. 1984. Variation among floral visitors in pollination ability: a precondition for mutualism specialization. – Science 225: 519–521.
- Stebbins, G. L. 1970. Adaptive radiation of reproductive characteristics in angiosperms, I: pollination mechanisms. – Annu. Rev. Ecol. Syst. 1: 307–326.
- Thomson, J. D. and Thomson, B. A. 1992. Pollen presentation and viability schedules in animal-pollinated plants: consequences for reproductive success. – In: Wyatt, R. (ed.), Ecology and evolution of plant reproduction. Chapman & Hall, pp. 1–24.
- Waser, N. M. 1979. Pollinator availability as a determinant of flowering time in ocotillo (*Fouquieria splendens*). – Oecologia 39: 107–121.
- Waser, N. M. and Campbell, D. R. In press. Adaptive speciation in flowering plants. – In: Dieckmann, U., Metz, H., Doebeli, M. and Tautz, D. (eds), Adaptive speciation. Cambridge Univ. Press.
- Waser, N. M., Chittka, L., Price, M. V. et al. 1996. Generalization in pollination systems, and why it matters. – Ecology 77: 1043–1060.

this Journal
is now available at

www.ingenta.com
the global research gateway

- full text PDF of all articles
- HTML abstracts and keyword search
- browse abstracts for free
- MEDLINE linking
- more than 2800 scientific journals



Need information about
Online Access? Please visit:

<http://journals.munksgaard.dk/ingentainfo>