— Note on Methodology —

USING THE LITERATURE TO TEST POLLINATION SYNDROMES — SOME METHODOLOGICAL CAUTIONS

Jeff Ollerton[1]*, André Rodrigo Rech[2], Nickolas M. Waser[3], and Mary V. Price[3]

[1]*Department of Environmental and Geographical Sciences, University of Northampton, Newton Building, Avenue Campus, Northampton, NN2 6JD, UK*
[2]*Institute of Biology, Federal University of Uberlândia, Av Pará, 1720 - Campus Umuarama. PO 593 - Uberlândia - 38400-902, Brazil*
[3]*School of Natural Resources and the Environment, University of Arizona, Tucson, AZ 85721, USA*

Abstract—"Pollination syndromes" are specific combinations of floral traits that are proposed to evolve convergently across angiosperm lineages in response to different types of animal pollinators. In spite of their long history, pollination syndromes have not been tested adequately–they rarely have been examined critically to determine how well they describe floral trait diversity or predict pollinators. In a recent meta-analysis of data from the literature, Rosas-Guerrero et al. (2014) provide a welcome test that draws on insights from past studies. At the same time, their study illustrates several difficulties of meta-analysis approaches in general, and for pollination biology in particular. Here we discuss those difficulties and propose some solutions. We first consider how to gather studies from the literature without introducing unintended bias, such as the old-fashioned method of working backward from cited literature. We next consider how to deal with difficulties that invariably arise when extracting and analyzing often-incomplete information from heterogeneous studies. Finally we discuss issues of interpreting and presenting the results in the most informative manner. We conclude that although Rosas-Guerrero et al. (2014) and other studies such as Ollerton et al. (2009) have arrived at different conclusions about the utility of pollination syndromes, their results are not necessarily incompatible.

Keywords: data interpretation, data presentation, literature searches, pollination syndromes, pollinator effectiveness, pollinator efficiency, meta-analysis

INTRODUCTION

The concept of pollination syndromes–suites of floral traits (colour, scent, morphology, etc.) that are associated with particular groups of pollinators–has played a central role in the field of pollination biology for almost one and a half centuries (Fægri & van der Pijl 1979; Waser 2006; Waser et al. 2011; Vogel 2012). Syndromes have identified apparent patterns in the bewildering diversity of floral phenotypes; they have been used to predict pollinators; and they have stimulated development of a core hypothesis in pollination ecology–that pollinator-mediated selection of floral traits has in large part driven the diversification of flowering plants. Various versions of pollination syndromes have been proposed, starting with those of Federico Delpino (1873-74; see also Waser et al. 2011). A version that is commonly used today is the set of trait descriptions elaborated by Faegri & van der Pijl (1979; hereafter the "traditional" syndromes).

Despite their importance, pollination syndromes have rarely been subjected to rigorous test. This is unsurprising

given how multi-faceted the syndromes are when viewed as a hypothesis. But it means that we do not know how well any particular version of the syndromes describes major patterns of variation in floral phenotypes or predicts plant-pollinator interactions. Such knowledge is critical for understanding the extent to which floral traits reflect adaptations to a "syndrome" pollinator rather than other possible adaptive scenarios (e.g. Waser 1998; Aigner 2001), and for devising improved schemes of floral classification.

One test of the traditional syndromes was attempted by Ollerton et al. (2009; hereafter "OLT"), who recorded visitation to and scored floral characteristics of 482 plant species in six communities from a variety of elevations, hemispheres, continents and latitudes. These authors asked (i) whether the traits of actual flowers, scored according to Faegri & van der Pijl's (1979) descriptions, formed clusters in a multivariate trait space that correspond with the locations of the idealized syndromes themselves (the answer was, "hardly ever"), and (ii) whether the most frequent pollinator observed at each plant species was predicted by the syndrome closest to it in phenotype space (the answer was, "about 30% of the time"). They concluded that the syndrome hypothesis ". . . as usually articulated does not successfully describe the diversity of floral phenotypes or predict the pollinators of most plant species".

Using a very different approach, Rosas-Guerrero et al. (2014; hereafter "R-G") have now provided a welcome additional test. R-G searched the existing literature to ask whether the traditional pollination syndromes correctly predict the "most effective" flower visitor of 417 plant species representing a variety of families and geographic regions. They concluded that syndrome pollinators are significantly more effective overall than non-syndrome pollinators, validating the traditional syndrome hypothesis.

The results of R-G are thought-provoking, and their meta-analysis approach holds great promise for drawing insights from the pooled time, energy, and experience of many researchers. At the same time, any single approach has both limitations and strengths for testing a hypothesis. Here we discuss how "the devil is in the details"—the utility of any meta-analysis with published literature depends on precisely how the literature is searched and analysed. We outline several difficulties and suggest some possible ways to deal with them. The difficulties have to do with *gathering representative studies* through a search of the literature; with *extracting and analyzing data* from heterogeneous sources; and with *interpreting and presenting the results* of the analysis. We conclude that the discrepant conclusions of R-G and OLT may be more apparent than real—the result of different methods and data—and we conclude by suggesting a strategy for continuing to improve our understanding of floral diversity.

GATHERING REPRESENTATIVE STUDIES

To evaluate the relationship between floral characteristics and the most effective pollinator, R-G searched the literature for data on syndromes and on pollinator effectiveness, a term we explain below. With modern web search engines, this may seem to be a straightforward task. But in reality, a search that is focused on just a few terms can turn up an incomplete and potentially unrepresentative sample.

In the example at hand, R-G searched the literature for variations on the terms "pollinator effectiveness", its synonym "pollinator efficiency", and "pollination syndromes". There are several potential limitations of this approach. First, a search using any term will miss studies published before the term came into common usage. "Pollinator effectiveness", for example, was rarely used before the 1970s (e.g. Stebbins 1970; Levin & Berube 1972; Primack & Silander 1975). Second, other relevant information appeared in papers in which the term "pollinator effectiveness" was not used, because data on effectiveness were incidental to the main topic. It may be impossible to say *a priori* how such limitations will affect the conclusions of any given meta-analysis, but it seems best to avoid the limitations as much as possible. In the case at hand several relevant studies that were missed by R-G immediately came to mind (e.g. Waser & Price 1981, 1990; Ollerton et al. 2003); below we will return to one of these.

There is a third limitation as well. The choice of search terms has the potential to capture an unrepresentative subset of studies even aside from the two concerns raised above. In this example, R-G's use of terms like "pollination

syndromes" may have led them disproportionately to studies where the authors were interested in the syndrome concept, rather than in some other topic, such as crop production. The effect of such unintended filtering could be a sample enriched in species that fit "interesting" syndromes.

What solutions do we propose? Our first recommendation (which reveals just how old some of the co-authors are!) is to augment an electronic search with the method in use before the web existed. This method is roughly as follows: identify recent papers on the topic at hand (one example for pollinator efficiency would be Ne'eman et al. 2010), identify relevant citations in them, read these earlier papers and identify relevant citations in them, and so on, working backwards as far as possible into the older literature. Such "backwards search" is time-consuming, but it can lead to information not revealed by a web search, including papers in a variety of languages, and thus a more complete overview of collective knowledge.

One must be careful that this strategy does not result in a narrow view of the literature caused by the tendency of authors to repeatedly cite the same publications, at the expense of more obscure literature. However, reading broadly in a field beyond the specific topic of interest (the agricultural literature, for example, is full of nuggets on pollination and plant reproduction), as well as directly contacting colleagues for their ideas of relevant and overlooked studies, combined with modern web searches, should result in a more comprehensive and representative assessment of the literature.

Our second recommendation is to avoid relying on any term in a web search that might lead to circularity (disproportionately finding studies that support a popular hypothesis), and instead to seek terms that are more neutral. In the case at hand this would mean augmenting variations on the term "syndromes" with terms such as "pollen deposition" and "pollinator visitation". Again, sifting through the large volume of resulting literature for papers that contain useful data is time-consuming, but likely to produce a more complete picture of what is known.

EXTRACTING AND ANALYSING THE DATA

a) Extracting syndrome data

The general issue here is how to extract information from the published studies that have been discovered in a literature search without introducing unintentional problems. This is an issue that will apply to any meta-analysis using previously-published studies, as several authors have pointed out (e.g. Gurevitch et al. 2001; Gates 2002). Here we focus on the specific example of pollination syndromes.

With studies of 417 plant species collated, R-G wished to determine whether the pollination syndromes of these species correctly predict the "most-effective" pollinators of flowers. The first step was to assign each plant species to a pollination syndrome. R-G expanded on the traditional syndrome descriptions of Faegri & van der Pijl (1979) to arrive at 11 diagnostic traits such as colour, shape, and reward type (their Tab. S1). Close inspection of their Tab.

S2 shows, however, that all 11 traits were scored for only 56 (13.4%) of the 417 flowers in R-G's sample. The remaining 361 flowers were scored only for various subsets of the traits (see below); in fact, one species was scored for only 5 of 11 traits, fewer than half. Such incomplete and variable scoring is an inevitable outcome of gathering studies from the literature, since there is no guarantee that all studies will measure a uniform set of traits. The consequence is that different sets of traits were used for different species, rather than all traits being employed with some standard weighting. Traits not used for a given species are effectively given weights of zero.

If plants are assigned to syndromes based on different sets of characters, then it is not clear what syndrome hypothesis is really being tested. The missing trait scores represent just that–missing data–so that syndrome assignments for different species are based on different amounts and kinds of information. If we examine the missing traits in R-G's data set, we see that virtually all species were scored for visual traits that are easy for humans to observe. These include flower colour, shape, orientation, size, symmetry, nectar guides, and position of anthers and stigmas. More species lacked information about traits that are harder for humans to observe: about 3% of species lacked reward type information, 12% lacked information on diurnal pattern of anthesis, 43% lacked odour strength information, and 79% lacked odour type information. Even accepting that odour strength and type are both highly subjective and generally poor descriptors of scent emissions (and therefore less frequently described by pollination ecologists), this suggests that pollination syndromes effectively were defined primarily by visual attributes, with a variable mix of other attributes playing a role only for subsets of species.

What can be done about missing data, which is likely to occur commonly in meta-analyses using existing studies? One idea is to try to fill in what is missing from other sources. In the current example this could include measuring herbarium specimens, consulting floras, or contacting knowledgeable people to determine missing floral traits. Once this is done as comprehensively as possible any patterns could then be re-analysed for the largest possible subset of species for which scores are available for the largest possible subset of traits, to see if results change. Another approach would be to begin with the present analyses and drop the trait that was most often missing, then the next most often missing, and so on, to see if results change. One benefit of these approaches could be to inform us of traits that add little to the predictive ability of pollination syndromes (also see Ollerton et al. 2009).

A related problem, which is not unique to literature meta-analysis approaches, is that investigators are likely to vary in how they perceive and score floral traits, and therefore how they assign species to syndromes. Ways of reducing such person-to-person variation are to have a single person score all species in the dataset, or to have multiple people do the scoring and arrive at a consensus syndrome assignment.

b) Estimating relative pollinator effectiveness

With the 417 plant species assigned to a syndrome, R-G next wished to compare the effectiveness of pollinators "expected" by the syndrome with the effectiveness of "unexpected" or "secondary" pollinators. They recognized that effectiveness has two multiplicative components, visit rate to flowers by pollinators (number of visits per unit of time) and pollen transfer per visit (see Stebbins 1970; Primack & Silander 1975). Unfortunately, many studies in the literature contain incomplete information about one or the other component of pollinator effectiveness. To get around this problem, some authors have argued that just one component provides an acceptable proxy for overall effectiveness: Vázquez et al. (2005) argued in favour of visit rate to flowers; King et al. (2013) felt strongly that pollen transfer per visit is the best proxy.

Using only one component as a proxy for overall pollinator effectiveness can be problematical for testing some aspects of the syndrome hypothesis unless the two components are positively correlated with one another–and available evidence suggests that often they are *not* correlated (e.g. Mayfield et al. 2001; Watts et al. 2012; Sakamoto & Morinaga 2013). Visit rate alone is fine if one is interested in how useful the syndrome concept is for predicting pollinators. But using only one component of pollinator effectiveness is insufficient for an overall evaluation of pollinator-mediated selection, which can involve floral traits that affect visit rate (pollinator attraction), as well as traits that affect per-visit pollen transfer.

R-G sided with King et al. (2013) and excluded studies that reported only pollinator visit rate. But then another sort of "missing data" problem presented itself–few studies directly measure per-visit pollen transfer. R-G therefore had to accept proxies for per-visit effectiveness. They used as proxies the amount of pollen removed from anthers or deposited on stigmas in a single visit, the amount of pollen on pollinators' bodies, contact with the flower's reproductive organs, or seed or fruit set.

Using a variety of proxies for pollinator effectiveness can complicate interpretation just as can using various subsets of floral traits to assign syndromes. A concrete example will illustrate this point. At one field site both bumblebees and hummingbirds visit *Delphinium nuttallianum* (= *nelsonii*), whose blue flowers and other traits strongly suggest the "bee syndrome". Indeed, each visit by a bee deposits on average about 10 times more pollen than each visit by a bird. But this component of effectiveness is *negatively* correlated with the other component: hummingbirds visit about 10 times more flowers per minute than do bees, and seed production attributable to the two functional groups is equivalent (Waser & Price 1981, 1990).

Interestingly, R-G included this species in their analysis but concluded that bumblebees are more effective based on a study (Waser 1978) that reported the difference in seed sets between uncaged plants that could be visited by both birds and bees, and caged plants that excluded birds. These two treatments averaged about 50 and 31 seeds per flower, respectively. However, it would be wrong to conclude that

each bee visit necessarily delivered 1.6 times more pollen than each bird visit. Instead, the treatment difference could result from differences in visit rate, if bees visited caged flowers more frequently than uncaged ones because the former accumulated more nectar in the absence of bird visits (Waser 1978, p. 939). Furthermore, seed and fruit set represent a post-pollination translation of stigma pollen loads (a function of both visit rate and per-visit deposition) into seeds or fruits, and this translation is nonlinear (see e.g., Waser & Price 1991). In this respect as well, seed and fruit set are not fool-proof indicators of either visit rate or per-visit deposition, or even of their product. To summarize, it is dangerous to use a difference in seed set between open-pollinated and partial exclusion treatments to judge the *quantitative* contributions of different pollinator functional groups. This subtle caution may apply as well to some of the other effectiveness calculations in R-G that are based on seed and fruit set, and it illustrates a more general issue in meta-analysis: the need for authors not only to assess quality of studies they use (as Gates 2002 emphasizes) but also to exercise caution in how they interpret results from those studies.

A useful way to assess the problem of heterogeneous data from literature meta-analysis is to conduct analyses on subsets of the data that are derived by different methods. R-G provide examples of this in their central analysis comparing effect sizes for the effectiveness of syndrome and non-syndrome pollinators (their Fig. 1b) for studies using different proxies for effectiveness or that differed in sample size.

c) Comparing pollinator effectiveness

R-G combine assessments of pollinator effectiveness and pollination syndrome to arrive at a measure for each plant species, Hedges' *d*, that reflects the difference in effectiveness of the syndrome pollinator vs. pollinators "not expected by the floral syndrome" (their p. 389 and Appendix 1). This is a very reasonable approach, but it is complicated by the fact that values for the effectiveness of "secondary" pollinators were available for only 57.5% of plant species, and many of these species had more than one "not expected" pollinator—another sort of "missing data" problem.

We do not see an easy way to include species with missing or multiple secondary pollinators without some sort of distortion. Assigning a value of zero for the effectiveness of the "not expected" pollinator when none was reported will automatically return a positive *d* value even if other pollinators do exist. In the case of multiple secondary pollinators, the value of *d* will similarly depend on whether it is based on the most-effective secondary pollinator, or the average effectiveness of all secondary pollinators.

Incomplete knowledge of ecological interactions bedevils any field study, because as Polis (1991) demonstrated, novel interactions continue to reveal themselves the longer one samples a single site, and the more sites one samples. Not only may new pollinator species and functional groups be identified as the sample expands in time and space, but some groups may be completely missing at specific times and places (e.g. Mayfield et al. 2001; Price et al. 2005; Alarcón

et al. 2008; Petanidou et al. 2008; also see discussion in Waser et al. 1996). Thus rankings of pollinator effectiveness might even depend on where and when a study was done.

Several of the species included in R-G illustrate this problem. To give just one example, R-G categorized *Daucus carota* (Apiaceae) as having flies as the syndrome pollinator group and wasps as the secondary pollinator group. However, *D. carota* is pollinated by various combinations of flies, wasps, and beetles in different years and sites (Lamborn & Ollerton 2000; Ollerton et al. 2007; Goulson et al. 2009).

What solutions do we suggest? First, we repeat the idea of doing analyses on subsets of the data that are derived by different methods. As mentioned above, R-G did do this (their Fig. 1b) by graphically presenting values of *d* based on different proxies for effectiveness. They also could have presented values of *d* separately for species with and without recorded secondary pollinators. Other than that, we can only suggest trying to characterize the degree to which data are incomplete and interpret results accordingly. For example, supplementary tables could list all known secondary pollinators, not just the ones for which quantitative data are available.

INTERPRETING AND PRESENTING THE RESULTS

In presenting and interpreting results of any study, it is most useful for the authors to clearly state the scope of the data, the assumptions necessary in analyses, and any other aspects that will help to place the results in context. Overlooking these details can lead to a mismatch between the results and how they are presented to readers.

One example in the present case is that R-G present their results as a "[test of the] universal character of the pollination syndromes". Any universal test would seem to require a sample representative of all the angiosperms, which is difficult if not impossible to achieve. In fact, the 417 species used by R-G belong to 76 different families (Tab. S2 of R-G, updated according to APGIII, 2009), only 18.7% of the 406 angiosperm families now recognized (www.theplantlist.org). Within this subset families are not represented according to their respective species numbers (Chi-square = 1760.4, df = 75, P < 0.00001) leading to under-representation of large families and over-representation of small families (Fig. 1). For example, Asteraceae (containing ca. 23,600 species) is represented by 5 species, Fabaceae (ca. 23,535 species) by 16, and Brassicaceae (ca. 3,505 species) by 2; whereas Aristolochiaceae (ca. 624 species) is represented by 6 species, Campanulaceae (ca. 2,385 species) by 15, Iridaceae (ca. 2,315 species) by 34, and Convolvulaceae (ca. 1,296 species) by 26 (24 in the single genus *Ipomoea*, whereas the family encompasses 67 genera). As a result of these limitations—yet another form of "missing data" —the species sampled by R-G are weighted toward large-flowered, relatively "showy" species with unusual, potentially specialised pollination systems (e.g. Sakai 2002; Muchhala & Potts 2007; Goldblatt et al. 1995; readers may contact the lead author for a compilation of images of 267 of the species in R-G's
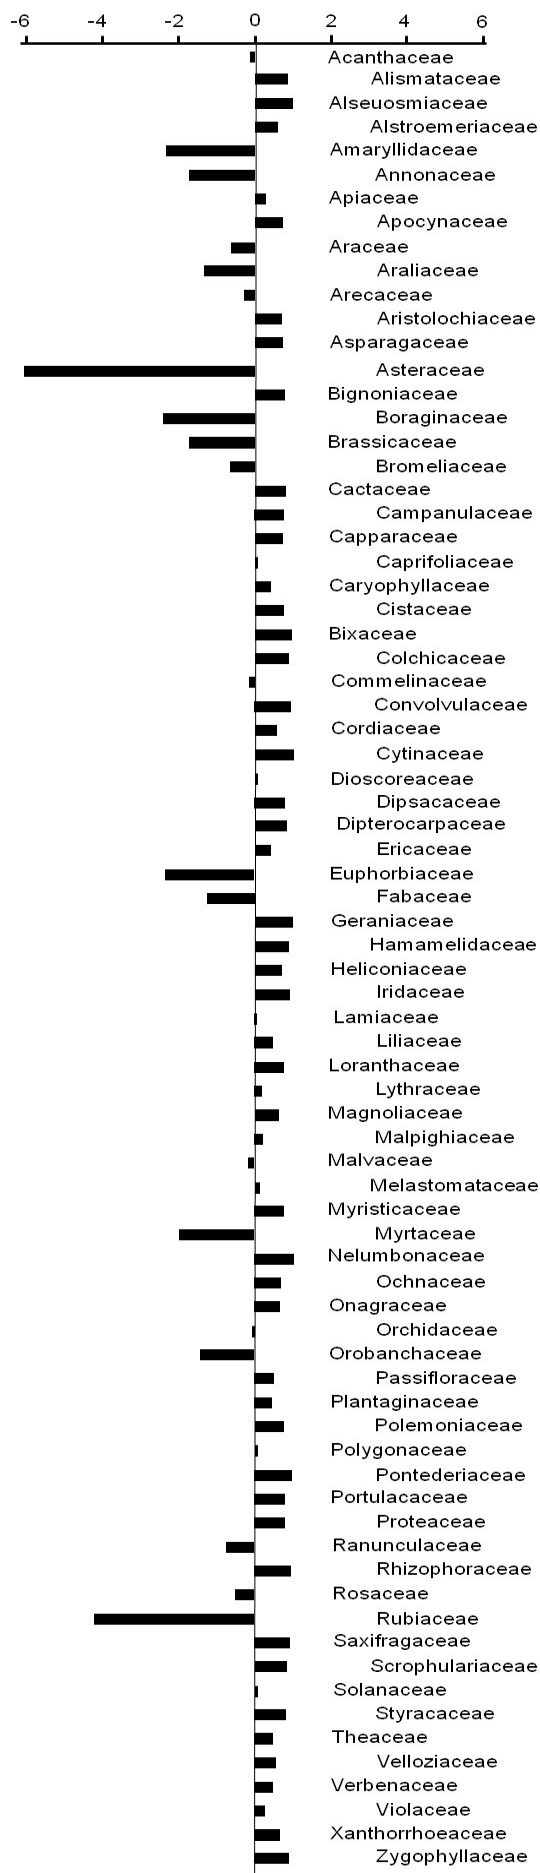
-6　　-4　　-2　　0　　2　　4　　6

Acanthaceae
　Alismataceae
Alseuosmiaceae
　Alstroemeriaceae
Amaryllidaceae
　Annonaceae
Apiaceae
　Apocynaceae
Araceae
　Araliaceae
Arecaceae
　Aristolochiaceae
Asparagaceae
　Asteraceae
Bignoniaceae
　Boraginaceae
Brassicaceae
　Bromeliaceae
Cactaceae
　Campanulaceae
Capparaceae
　Caprifoliaceae
Caryophyllaceae
　Cistaceae
Bixaceae
　Colchicaceae
Commelinaceae
　Convolvulaceae
Cordiaceae
　Cytinaceae
Dioscoreaceae
　Dipsacaceae
Dipterocarpaceae
　Ericaceae
Euphorbiaceae
　Fabaceae
Geraniaceae
　Hamamelidaceae
Heliconiaceae
　Iridaceae
Lamiaceae
　Liliaceae
Loranthaceae
　Lythraceae
Magnoliaceae
　Malpighiaceae
Malvaceae
　Melastomataceae
Myristicaceae
　Myrtaceae
Nelumbonaceae
　Ochnaceae
Onagraceae
　Orchidaceae
Orobanchaceae
　Passifloraceae
Plantaginaceae
　Polemoniaceae
Polygonaceae
　Pontederiaceae
Portulacaceae
　Proteaceae
Ranunculaceae
　Rhizophoraceae
Rosaceae
　Rubiaceae
Saxifragaceae
　Scrophulariaceae
Solanaceae
　Styracaceae
Theaceae
　Velloziaceae
Verbenaceae
　Violaceae
Xanthorrhoeaceae
　Zygophyllaceae

FIGURE I. Taxonomically-biased representation of plant families in the R-G analysis. The bars show I − [(proportional representation in R-G)/(proportional representation among all angiosperm species)] for each family. A value of zero indicates that the family is sampled in relation to its size, a negative value that it is under-represented in R-G, and a positive value that it is over-represented.

dataset), rather than representing a universal picture of the flowering plants as a whole. In short, the literature used by R-G (and perhaps the literature in general) is biased towards studies of rather specialised taxa and against studies of more generalised plants.

We wish to turn finally to an issue of data presentation that occurred to us as we read R-G and looked at the figures. This is a realization about the graphical representation of results of a meta-analysis and how easily it can be unintentionally misleading. Fig. Ia of R-G shows mean Hedges' d, and 95% confidence interval of the mean, for various subsets of species and for all data–a standard way of reporting results of meta-analysis. Technically speaking, what a 95% confidence interval of the mean indicates (under certain assumptions) is the range of values that will include the true (unknown) population mean in 95% of repeated samples of the same size from the same underlying population. As such, this statistic is appropriate for showing whether a mean value exceeds zero, which is how it is used in R-G and in other meta-analyses. But the visual effect of Fig. Ia of R-G is to suggest that the actual d values cluster tightly around a very positive value, and in fact this is not the case. Instead, Tab. S2 of R-G reveals that d values were zero or negative for I02 of 4I7 (24.5%) of the plant species, indicating equal or even *higher* effectiveness of "secondary pollinators" relative to "syndrome pollinators". This variation among species would be better illustrated by showing $\pm$ I.96 standard deviations of the actual d values, i.e., the interval within which (at least approximately) 95% of the values for individual species lie. Our Fig. 2 contrasts these two ways of presenting the core results of R-G's study. Whereas the 95% confidence interval of the overall sample mean d of ca. 0.65 ranges from ca. 0.6 to 0.7, far from the negative territory that contains 24.5% of all d values, the standard deviation in these values is I.077, so that $\pm$ I.96 standard deviations range from + 2.76 to - I.46, well within negative territory.

A related point is that d values that are positive but of small magnitude may not indicate any important superiority of the "syndrome" pollinator. For example, consider that 92 d values in Tab. S3 of R-G are positive but less than 0.6. This is approximately the lower confidence limit in Fig. Ia of R-G and represents a case in which the primary pollinator is better than the secondary one by about six tenths of a weighted standard deviation in the effectiveness of the two pollinators–not an overwhelming difference. Compared to a maximum d value of 3.68, then, 46.5% or nearly half of all values were negative, zero, or positive but of small magnitude. It is questionable whether small positive values can be taken to mean that the primary pollinators were so much more effective than the secondary ones that they were
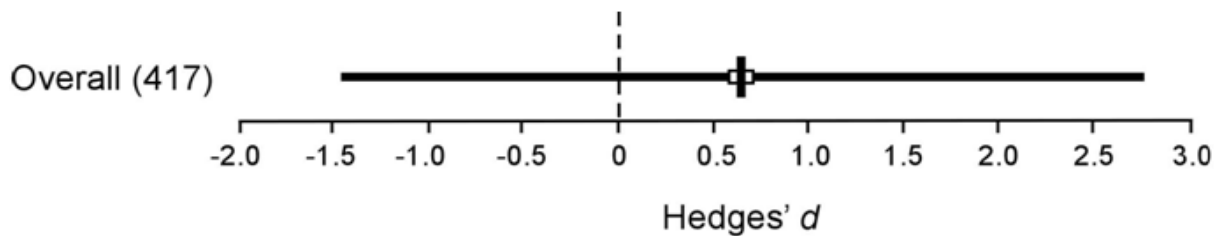
FIGURE 2.　Approximately 95% of the 417 values of Hedges' *d* reported in R-G's Tab. S2 fall within the range shown by the dark horizontal bars around the mean value of 0.65, shown by the vertical bar. By way of contrast, the small rectangle around this mean value represents the 95% confidence interval of the mean as displayed in Fig. 1a of R-G.

exclusively selecting for those flowers' phenotypes in a strictly adaptive explanation of the syndromes.

COMPARISON WITH OLT AND A WAY FORWARD

OLT evaluated the traditional syndromes by asking how often the most frequent pollinator of the 483 plant species observed in six communities by similar methods was the one predicted by the closest syndrome. In contrast, R-G asked whether the average difference in effectiveness of expected (= syndrome) and unexpected pollinators reported for 417 species in the literature is positive. Comparing the two studies by these measures, the syndromes made the "correct" prediction for about 30% of the species in OLT's sample and about 75% in R-G's sample (if one accepts even small positive values of Hedges' *d*). But if we consider possible effects of differences between these two studies in the mix of species used, methods for assigning plants to syndromes, measures of pollinator effectiveness, and analysis, the true discrepancy in results may be smaller or even zero.

It may be a human tendency to dichotomize issues (see Bertens 2013 on "binary opposition"), or in the specific example at hand to assume that there are two possible philosophical positions: that pollination syndromes either are baseless or conversely that they are universally valid. But the results of both OLT and R-G show that this dichotomy is false and that the applicability of traditional syndromes falls somewhere between about 30% and 75% of angiosperm species. Therefore, we do not perceive that there is a "strong debate about the reliability of pollination syndromes", as R-G suggest, but instead a less-polarised scientific landscape in which their conclusions and OLT are not so much at odds, as they are based on different methods and data. This is backed up by recent studies suggesting, for instance, that quantitative (as opposed to categorical) trait descriptions may be more reliable for predicting the main pollinators of a plant, especially if they incorporate a more nuanced view of the relative importance of different sensory modalities for each pollinator taxon (Junker et al. 2013; Junker & Parachnowitsch 2015).

Our hope is that this conceptual landscape can be further explored carefully and dispassionately in the future. In this task we are firmly in favour of more meta-analyses to extract useful information from many decades of published studies. Such analyses can be a much swifter way to accumulate information than the very time-consuming *de novo* empirical surveys such as those of OLT (which required effort over

almost an entire decade by a team of workers). In conjunction with literature surveys, however, further empirical sampling will continue to be an essential source of new data that will illuminate natural patterns that may be hard to detect from analyses of the older literature. As part of this future empirical work, it will be most useful to record both expected–and unexpected–trait expressions of flowers and identities of flower visitors. Such an open-minded approach will nurture the continued evolution of our understanding of the patterns that led early workers to codify the pollination syndromes, by improving our assessment of how floral traits relate to visitors, of which traits are most informative, of the ways in which different traits represent adaptations, and of whether patterns are unchanging across taxa and geographic regions.

REFERENCES

Aigner, PA (2001). Optimality modeling and fitness tradeoffs: when should plants become pollinator specialists? Oikos 95:177–184.

Alarcón R, Waser NM, Ollerton J (2008) Annual variation in the topology of a plant-pollinator interaction network. Oikos 117:1796–1807.

A.P.G. III (2009). An update of the Angiosperm Phylogeny Group classification for the orders and families of flowering plants. Botanical Journal of the Linnean Society 161:105–121.

Bertens H (2013) Literary Theory: the Basics. Routledge, New York.

Delpino F (1873–1874) Ulteriori osservazioni e considerazioni sulla dicogamia nel regno vegetale. Atti della Società Italiana di Scienze Naturali 16:151–349, 17:266–407.

Faegri K, van der Pijl L (1979) The Principles of Pollination Ecology, 3rd Revised Edition. Pergamon Press, Oxford .

Gates, S (2002) Review of methodology of quantitative reviews using meta-analysis in ecology. Journal of Animal Ecology 71:547–557.

Goldblatt P, Manning JC, Bernhardt P (1995) Pollination biology of *Lapeirousia* Subgenus *Lapeirousia* (Iridaceae) in Southern Africa; floral divergence and adaptation for long-tongued fly pollination. Annals of the Missouri Botanical Garden 82:517–534.

Goulson D, Mcguire K, Munro EE, Adamson S, Colliar L, Park KJ, Tinsley MC, Gilburn AS (2009) Functional significance of the dark central floret of *Daucus carota* (Apiaceae) L.; is it an insect mimic? Plant Species Biology 24:77–82.

Gurevitch, J, Curtis PS, Jones MH (2001) Meta-analysis in ecology. Advances in Ecological Research 32:199–247.

Hingston AB, McQuillan PB (2000) Are pollination syndromes useful predictors of floral visitors in Tasmania? Austral Ecology 25: 600–609.

Junker RR, Blüthgen N, Brehm T, Binkenstein J, Paulus J, Schaefer HM, Stang M (2013) Specialization on traits as basis for the niche-breadth of flower visitors and as structuring mechanism of ecological networks. Functional Ecology 27:329–341.

Junker RR, Parachnowitsch AL (2015) Working towards a holistic view on flower traits—how floral scents mediate plant–animal interactions in concert with other floral characters. Journal of the Indian Institute of Science 95:43–67.

King C, Ballantyne G, Willmer PG (2013) Why flower visitation is a poor proxy for pollination: measuring single-visit pollen deposition, with implications for pollination networks and conservation. Methods in Ecology and Evolution 4: 811–818.

Lamborn E, Ollerton J (2000) Experimental assessment of the functional morphology of inflorescences of *Daucus carota* (Apiaceae): testing the "fly catcher effect". Functional Ecology 14:445–454.

Levin DA, Berube DE (1972) *Phlox* and *Colias*: the efficiency of a pollination system. Evolution 26: 242–250..

Mayfield MM, Waser NM, Price MV (2001) Exploring the "most effective pollinator principle" with complex flowers: bumble bees and *Ipomopsis aggregata*. Annals of Botany 88:591–596.

Muchhala N, Potts MD (2007) Character displacement among bat-pollinated flowers of the genus *Burmeistera*: analysis of mechanism, process and pattern. Proceedings of the Royal Society B 274:2731–2737.

Ne'eman, G, Jürgens A, Newstrom-Lloyd L, Potts SG, Dafni A (2010) A framework for comparing pollinator performance: effectiveness and efficiency. Biological Reviews 85: 435–451.

Ollerton J, Johnson SD, Cranmer L, Kellie S (2003) The pollination ecology of an assemblage of grassland asclepiads in South Africa. Annals of Botany 92:807–834.

Ollerton J, Killick A, Lamborn E, Watts S, Whiston M (2007) Multiple meanings and modes: on the many ways to be a generalist flower. Taxon 56:717–728.

Ollerton J, Alarcón R, Waser NM, Price MV, Watts S, Cranmer L, Hingston A, Peter CI, Rotenberry J (2009) A global test of the pollination syndrome hypothesis. Annals of Botany 103:1471–1480.

Pauw A (2006) Floral syndromes accurately predict pollination by a specialized oil-collecting bee (*Rediviva peringueyi*, Melittidae) in a guild of South African orchids (Coryciinae). American Journal of Botany 93:917–926.

Petanidou T, Kallimanis AS, Tzanopoulos J, Sgardelis SP, Pantis JD (2008) Long-term observation of a pollination network: fluctuation in species and interactions, relative invariance of network structure and implications for estimates of specialization. Ecology Letters 11:564–575.

Polis GA (1991) Complex trophic interactions in deserts: an empirical critique of food-web theory. American Naturalist 138:123–155.

Price MV, Waser, NM, Irwin RE, Campbell DR, Brody AK (2005) Temporal and spatial variation in pollination of a montane herb: a seven-year study. Ecology 86:2106–2116.

Primack RB, Silander JA (1975) Measuring the relative importance of different pollinators to plants. Nature 255:143–144.

Rosas-Guerrero V, Aguilar R, Marten-Rodriguez S, Ashworth L, Lopezaraiza-Mikel M, Bastida JM, Quesada M (2014) A quantitative review of pollination syndromes: do floral traits predict effective pollinators? Ecology Letters 17: 388–400.

Sakai S (2002) *Aristolochia* spp. (Aristolochiaceae) pollinated by flies breeding on decomposing flowers in Panama. American Journal of Botany 89:527–34.

Sakamoto RL, Morinaga S (2013) Poor correlation between the removal or deposition of pollen grains and frequency of pollinator contact with sex organs. Naturwissenschaften 100:871–876.

Stebbins GL (1970) Adaptive radiation of reproductive characteristics in angiosperms. I: pollination mechanisms. Annual Review of Ecology and Systematics 1:307–326.

Vázquez DP, Morris WF, Jordano P (2005) Interaction frequency as a surrogate for the total effect of animal mutualists on plants. Ecology Letters 8:1088–1094.

Vogel S (2012) Floral-biological syndromes as elements of diversity within tribes in the flora of South Africa. Shaker Verlag, Aachen.

Waser NM (1978) Competition for hummingbird pollination and sequential flowering in two Colorado wildflowers. Ecology 59:934–944.

Waser, NM (1998) Pollination, angiosperm speciation, and the nature of species boundaries. Oikos 82:198–201.

Waser NM (2006) Specialization and generalization in plant-pollinator interactions: an historical perspective. In Waser NM, Ollerton J (eds) Plant-Pollinator Interactions: from Specialization to Generalization. University of Chicago Press, Chicago, pp 3–17.

Waser NM, Chittka L, Price MV, Williams N, Ollerton J (1996) Generalization in pollination systems, and why it matters. Ecology 77:1043–1060.

Waser NM, Ollerton J, Erhardt A (2011) Typology in pollination biology: lessons from an historical critique. Journal of Pollination Ecology 3:1–7.

Waser NM, Price MV (1981) Pollinator choice and stabilizing selection for flower color in *Delphinium nelsonii*. Evolution 35:376–390.

Waser NM, Price MV (1990) Pollination efficiency and effectiveness of bumble bees and hummingbirds visiting *Delphinium nelsonii*. Collectanea Botanica (Barcelona) 19:9–20.

Waser NM, Price MV (1991) Outcrossing distance effects in *Delphinium nelsonii*: pollen loads, pollen tubes, and seed set. Ecology 72:171–179.

Watts S, Huamán Ovalle D, Moreno Herrera M, Ollerton J (2012) Pollinator effectiveness of native and non-native flower visitors to an apparently generalist Andean shrub, *Duranta mandonii* (Verbenaceae). Plant Species Biology 27:147–158.