GM Crops 1:1, 25-31; January/February 2010; © 2010 Landes Bioscience

Reducing uncertainty in regulatory decision-making for transgenic crops

More ecological research or clearer environmental risk assessment?

Alan Raybould

Syngenta; Jealott's Hill International Research Centre; Bracknell, Berkshire UK

Key words: hypothesis testing, risk hypothesis, regulatory risk assessment, environmental policy

Ecological research and environmental risk assessment are similar in that they address interesting problems by formulating and testing hypotheses. They differ in the types of problems that are interesting, the characteristics of good hypotheses to solve those problems, and the methods for rigorous testing of hypotheses. It is important to recognize the differences between environmental risk assessment and basic ecological research because confusing them can lead to ineffective risk assessment and missed opportunities to advance ecological theory. Uncertainty in regulatory decision-making about transgenic crops may be reduced more effectively by clarifying the purpose and structure of environmental risk assessments than by further research on the ecology of the crops.

Introduction

Uncertainty about environmental risks is often cited as the reason why decision-making about approvals for commercial cultivation of transgenic crops can be protracted and contentious.[1] More research may appear the obvious solution to increasing confidence in risk assessments and thereby quickening and improving trust in decision-making. One problem with this assumption is that accumulating research suggesting that the cultivation of transgenic crops poses low environmental risk and provides many environmental benefits[2,3] has not resulted in shorter decision times or reduced controversy over applications for cultivation of transgenic crops, particularly in Europe where much of the research has been done.[4]

Craig et al. recognized the potential for environmental risk assessment research to overwhelm regulators, and thereby increase the time taken to evaluate regulatory dossiers: "In the decade since the first authorizations for commercial release of transgenic crops, there has been an enormous increase in the amount of data generated by scientific studies that relates to risk assessment. If this trend continues, we run the risk of competent authorities being submerged by excessively large amounts of data that may be of questionable pertinence to verifiable safety questions."[5] Time spent evaluating the burgeoning literature may not be the only factor contributing to lengthy decision times. As Johnson et al. pointed out, "collecting data and making vague assertions that they are relevant to risk assessment, without providing specific predictions about things of concern, only serves to confuse and increase unease."[6] Thus irrelevant data (i.e. "data that may be of questionable pertinence to verifiable safety questions") may affect decision-making not only because they take time to review, but also because they reduce confidence that a decision is correct. Reduced confidence, or increased unease and uncertainty about decisions may lead to a cycle that makes the problem worse: the public does not trust regulators and the plant biotechnology industry; regulators of transgenic crops must appear to be thorough to try to increase public trust; regulators commission more research; the research increases confusion, reduces confidence in decisions and reduces public trust, and so on.

This paper examines the similarities and differences between risk assessment and basic research. It first examines two misunderstandings about risk assessment that may increase the amount of research unhelpful to risk assessment. Once these misunderstandings are recognized, risk assessment can be seen to have the same logical structure as basic research: the identification of interesting problems and the formulation and testing of hypotheses as possible solutions to those problems. The paper then discusses important differences in the selection of problems, the criteria for finding hypotheses interesting and the methods for testing hypotheses. Failure to recognize these differences harms both risk assessment and basic research: approaching risk assessment as a basic research problem often does not produce information useful for decision-makers; and attempting to make basic research relevant to risk assessment distracts scientists from important theoretical questions, and means that opportunities to advance fundamental knowledge are missed. The paper concludes that a clearer demarcation between environmental risk assessment

*Correspondence to: Alan Raybould; Email: alan.raybould@syngenta.com
Submitted: 07/15/09; Accepted: 08/11/09
Previously published online:
www.landesbioscience.com/journals/gmcrops/article/9776

and basic research on the ecology of transgenic crops should improve both activities.

Two Misunderstandings about Risk Assessment

The correct policy can be discovered by research. Risk assessment estimates the likelihood and seriousness of harm that may result from the cultivation of a transgenic crop.[7] A risk assessment does not conclude whether or not the transgenic crop should be cultivated; it is one source of information that helps decision-makers to judge whether cultivating the crop will achieve objectives of policy.[6,8]

A common approach to applying science to policy is the "deficit model," which supposes that poor policy results from a deficit of knowledge.[9] The deficit model implies that policy can be formulated directly from research, or even that certain policies are inevitable given particular discoveries. The problem with the deficit model is that policy must reconcile different preferences and settle competing demands, not ensure or prevent certain results at all costs. This was elegantly put by Lubechenko:[10] "Many of the choices facing society are moral and ethical ones, and scientific information can inform them. Science does not provide the solutions, but it can help understand the consequences of different choices." Thus, scientific research clarifies the implications of taking particular actions, it alone cannot decide whether those implications are desirable and which actions should be taken.

It follows that in risk assessment for transgenic crops, what constitute harmful effects are decided by subjective judgements and need to be defined by policy prior to a risk assessment; a risk assessment is not research to discover what should be regarded as harmful. Research cannot substitute for clear definitions of harm in laws or regulations; however, the objectives of much ecological research related to transgenic crops, such as the Farm Scale Evaluations in the UK,[11] suggest that it is frequently done in an attempt to fill a gap that should be occupied by policy. Numerous field studies have sought to compare as many variables as possible between fields of transgenic and non-transgenic crops. This method implies that there is no clear definition of harm, and that harm is defined post hoc as changes detected by the experiment. This approach is inefficient, because many variables of no value will be measured; it is ineffective, because experiments designed specifically to test whether harm will arise would be more informative to decision-makers; and it results in capricious decisions, because what changes will differ between experiments.[12]

Scientists, regulators and decision-makers must avoid the Naturalistic Fallacy[13] of thinking that because basic research tells us how the world *is*, it can also tell us how the world *ought to be*. If the problem facing decision-makers is uncertainty about what is valuable and what is harmful, scientific research will not solve that problem, and it may make it worse.

Risk assessments prove safety. It is a common misconception that risk assessment is a search for proof that an activity is safe. This is a specific case of the widely held, or at least often stated, idea that science proves that certain theories are true by collecting observations that support them. In other words, science draws general conclusions from particular cases, a method called induction. The many problems of induction as a logical basis for science are well known.[14] They include the impossibility of proving that no observation will be made that contradicts a theory, and that many theories may fit any set of observations. The classic example to demonstrate problems of induction is to assert that if one observes only white swans, eventually the theory that all swans are white will be proven; however, observing any number of white swans cannot eliminate the possibility that a black (i.e., non-white) swan will be observed, and many other theories about swan coloration would fit one's observation that every swan is white. Nevertheless, despite the logical difficulties, science may appear to advance by induction, but this is a problem of psychology[15] or possibly of semantics.[16]

Thinking about swans clarifies why safety cannot be proven. If a harmless result of performing an activity is equivalent to a white swan, and a harmful result is equivalent to a black swan, then just as one cannot prove that all swans are white, one cannot prove that every time the activity is performed no harm will result; in other words, it cannot be proved that the activity is safe. Attempts to prove safety are futile, and may result in the production of data in excess of what is needed for sound decision-making.

Similarities between Risk Assessment and Research: Problem Solving by Hypothesis Testing

If safety cannot be proved by induction, how is it possible to demonstrate that a risk is acceptable? Popper offered a solution to the "problem of induction" by proposing that science proceeds by deduction; that is, it infers particular facts from general propositions. In the case of swans, one may have a theory that all swans are white, and from that theory one would deduce that all the swans on a particular river at a particular time will be white. If indeed one observed that all the swans on that river were white, the observations would *corroborate* the theory that all swans are white, but the theory would not be proved, no matter how many times one observed white rather than non-white swans.

The difference between induction and deduction may appear trivial for theories about swan coloration; however, the difference has important implications. Popper argued that science proceeds by the identification of problems and the formulation of trial or tentative solutions to those problems; the trial solutions are then tested to eliminate error, which leads to new knowledge and new problems.

A simple scheme represents this idea[15]

$$P_1 \square TS_1 \square EE_1 \square P_2 \square$$

where P_1 is an initial problem or question, which arises from attempts to solve prior problems, and could be a set of observations that contradict a previous hypothesis; TS_1 is a trial or tentative solution to the problem, such as a new hypothesis that explains the observations; EE_1 is error elimination, attempts to remove flaws in the hypothesis by putting it to the test; and P_2 is new knowledge and a new problem or set of problems for which new trial solutions may be proposed.

The crucial point is that science begins with problems that lead to testable hypotheses and attempts to *falsify* those hypotheses;

science does not observe the world without preconceptions (i.e., theories) and create theories from those observations, and then seek further observations to *support* the theory with a view to proving it. In the case of swans, knowledge of swan coloration is increased by searching for non-white swans to test the hypothesis that all swans are white. The theory is put under the most severe tests by observing swans in places where non-white swans are most likely to be found; if no non-white swans are found, the theory has survived rigorous testing and its corroboration is strengthened. If black swans are observed, then the all-white swan theory is falsified and a new theory is needed. Knowledge of swan coloration is not gained simply by observing a few white swans and then searching for more white swans to prove that all swans are white. Observations are always made with some purpose or expectation (one cannot simply "observe"[17]) and counter-examples to one's theories should be actively sought to eliminate errors and thereby gain knowledge.

Basic research can therefore be viewed as the testing of hypotheses that are tentative solutions to interesting scientific problems. Scientific knowledge is a collection of explanatory theories that have withstood falsification under test. Knowledge advances as further tests reveal the errors in those theories, leading to new theories and further testing. Risk assessment of the cultivation of transgenic crops also fits this model.[18]

The initial problem is the identification of potential harmful effects that may arise from cultivation of the transgenic crop. The trial solution to that problem is one or more hypotheses that cultivation of the crop will not lead to the identified harm; these are known as risk hypotheses. The risk hypotheses are tested using existing or new studies designed to show that cultivation of the transgenic crop will be harmful, in other words to try to falsify the risk hypotheses. If no harmful effects are observed, the hypotheses have survived rigorous testing, and the risks posed by cultivation of the transgenic crop can be regarded as low.[12]

To summarize, acceptable risk is a judgment that risk hypotheses predicting no harm from the cultivation of a transgenic crop have been sufficiently corroborated. This is analogous to a judgment in basic research that a particular hypothesis has been sufficiently corroborated and that knowledge would be advanced more fruitfully by investigation of a new problem with its different associated hypotheses. What are regarded as harmful effects of a transgenic crop must be defined, they cannot be discovered by the risk assessment, just as tests of theories about swan coloration say nothing about which color of swan should be preferred. The risk hypotheses are corroborated not by seeking *evidence for* no harmful effects, but by *testing* hypotheses of no harmful effects: one searches for harmful black swans, not for harmless white ones. Finally, none of the above means that hypotheses should never be re-examined, nor that the acceptability of a risk should be never be re-evaluated, if new information makes that worthwhile.

Differences between Risk Assessment and Basic Research

Recognizing that risk assessment follows the same deductive logic as basic research is important; nevertheless, risk assessment and basic research differ in the sources of problems, the nature of the hypotheses under test, and even the methods for testing hypotheses. Unless these differences are recognized, risk assessment and basic research become confused to the detriment of both activities: risk assessment fails to provide information of immediate use to decision-makers; basic research fails to test theories that advance fundamental knowledge and which may provide valuable new methods for future risk assessments.

Problem selection. In Popper's model of the development of objective knowledge through deduction, the first stage is identification of an interesting problem. In science, the source of new problems is the testing of hypotheses that are trial solutions to older problems. A hypothesis is tested by objectively comparing its predictions with observations made under conditions that have the potential to reveal its flaws. Objective comparison of predictions and results may produce new problems, which may give the impression that subjectivity has no place in basic research; nevertheless, selection of which new basic research problems to try to solve does contain subjective elements because of personal preferences of the scientist[17,19] or societal interests, perhaps mediated through the allocation of research grants.[20]

A belief in the objectivity of problem selection may not harm basic research; it is, however, a major problem for risk assessment. Problem selection for risk assessment must be subjective because its purpose is to help protect things we value from being harmed—value and harm being subjective criteria. Attempts at objectivity in the selection of problems for risk assessment may make scientists avoid basing their work on definitions of harm so as not to appear biased; thus, the problem to be solved becomes "what will happen when this transgenic crop is cultivated?" not "what is the probability that cultivation of the transgenic crop will cause harm?"[18] This approach can lead to exhaustive comparisons of fauna between fields of transgenic insect resistant[21] or transgenic herbicide-tolerant crops and suitable non-transgenic comparators,[11] or to complex mathematical models that investigate the logical consequences of assumptions about how the transgenic crop interacts with the environment.[22] At best, these "objective" approaches are inefficient, and at worst they risk the Naturalistic Fallacy: in this case, equating any change associated with the transgenic crop, or any differences between the transgenic and non-transgenic crops, as harmful effects. Effective risk assessment must select problems subjectively based on prior definitions of harm, not conduct research to describe possible changes or catalog differences.

Formulation of hypotheses. Scientific hypotheses seek to be accurate solutions to interesting problems. While hypotheses may be logical deductions from existing knowledge, hypothesis formulation requires imagination from the scientist to devise and select a particular trial solution from the many alternatives that may be tenable based on existing knowledge. The selection of a particular hypothesis as the best trial solution can be seen as making assumptions; however, these assumptions are not guesses, but plausible deductions from existing knowledge.

An important difference between basic research and risk assessment is in what is considered interesting. In basic research, an interesting hypothesis is "a bold anticipation of things to come";[14] it has a high degree of improbability and testability, which often

Table 1. A general scheme for formulating risk hypotheses

Scenario	Hypothesis
Cultivation of the GM crop	
□	Event A will not occur
Event A	
□	Event B will not occur
Event B	
□	Event C will not occur
Event C	
□	Event D will not occur
Event D (Harm)	

A scenario is a pathway by which harm may arise from cultivation of the GM crop. A risk hypothesis postulates that a particular step in the pathway is absent or occurs at a frequency or magnitude below a threshold.

Table 2. Derivation of risk hypotheses to characterize the risk from gene flow from TuMV-resistant oilseed rape to wild species

Scenario	Risk hypothesis
Cultivation of transgenic TuMV resistant oilseed rape	
□	No hybridization between oilseed rape and the wild species
Hybridization between the transgenic oilseed rape and a wild species	
□	The wild species is immune to TuMV
The transgene increases the TuMV resistance of the wild species	
□	TuMV does not infect the wild species in the field
The wild species is infected by TuMV in the field	
□	TuMV infection does not reduce seed production
Infected transgenic plants produce more seed than infected non-transgenic plants	
□	Abundance of the wild species is not limited by seed production
Increased abundance of the wild species which reduces the abundance of values species (harm)	

means that it makes precise predictions. It is easy to make accurate predictions: the hypothesis that it will rain somewhere in Europe in the next month is likely to be accurate, but it is hardly a bold anticipation—it is uninteresting because it is so likely to be corroborated. On the other hand, a hypothesis that it will be raining in Madrid, but not in London, at 3 pm next Friday, is less likely to be accurate. The precision of its predictions make the hypothesis more interesting because of the greater explanatory power of the underlying theory.

In risk assessment, the interest of the hypothesis comes from its ability to improve decision-making; thus, if an accurate prediction of whether or not it will rain in Europe in the next month is sufficient to make a decision, building and testing hypotheses about precisely where and when it will rain is unnecessary. In fact, testing a hypothesis that makes very precise predictions may confuse decision-making as it implies that the exact timing and location of rainfall should be considered even though a decision can be made without that information.

An example of the difference between an interesting research hypothesis and a risk hypothesis is the work of Wilkinson et al.[23] on hybridization between oilseed rape (*Brassica napus*) and wild turnip (*B. rapa*) in the United Kingdom. Environmental harm may arise from the introgression of a transgene from a transgenic crop into a wild species if the transgene leads to "ecological release," that is the wild species becomes protected from an environmental factor that was limiting its abundance, and increased abundance of the species causes harm.[24] Wilkinson et al.[23] used a combination of remote sensing, field work and mathematical modeling to predict the number of oilseed rape x wild turnip hybrids produced in the United Kingdom each year. The spatial distribution of the hybrids was also estimated. It is not clear how a precise prediction of the number of hybrids helps to assess the risks from cultivation of transgenic oilseed rape. Previous work[25] had shown that oilseed rape and wild turnip hybridize spontaneously in the field, and the risk from a transgenic hybrid depends on the potential for the transgene to cause ecological release. While a hypothesis that, say, 30,000 hybrids will form annually is an interesting hypothesis for basic research, it is unnecessarily precise for risk assessment. A risk hypothesis that oilseed rape and wild turnip do not form hybrids in the field is sufficient (see below), and this hypothesis had already been falsified,[25] suggesting that risk assessments for cultivation of transgenic oilseed rape in the United Kingdom should concentrate on the properties of the transgenes, not on increasing the precision of estimates of the number of hybrids.

In general, risk hypotheses are formulated by considering pathways or scenarios by which the activity under evaluation, in this case the cultivation of a transgenic crop, will cause harm. A specific risk hypothesis postulates a break in the chain of events of a particular scenario (**Table 1**), which may be the complete absence of a phenomenon, or that an event occurs below a certain frequency or threshold magnitude. As described above, the derivation and selection of risk hypotheses for testing in the risk assessment could be seen as "making assumptions," and unnecessarily restricting the risk assessment; however, many pathways to harm may be so implausible—that is the risk hypotheses that arise from them are sufficiently corroborated with existing data—that they do not require evaluation in the risk assessment. It may be necessary for good risk communication to state that particular pathways were considered and rejected as implausible,[7] but it is not necessary to test every possible risk hypothesis, and certainly not with new studies.

A real scenario for harm that may result from cultivation of a transgenic crop is described by Raybould and Cooper.[24] They investigated the environmental risks posed by the cultivation of turnip mosaic virus (TuMV) resistant transgenic oilseed rape in the United Kingdom, and in particular a scenario whereby the virus resistance transgenes introgress into populations of wild

Table 3. A comparison of possible risk hypotheses and research hypotheses derived from particular scenarios whereby a transgenic crop could cause harm

Scenario	Risk hypothesis	Research hypothesis
Transgenic produces a protein toxic to an endangered species	NOAEC > 10X highest exposure	NOAEC = X ug/g diet Exposure = Y ug/g diet
Introgression of a transgene causes increase in abundance of a wild species	Crop A does not hybridize with wild plant B	Number of A x B hybrids = X
Transformation creates a noxious weed	Transgenic crop A is no more weedy than its progenitor B	Number of weeds of transgenic crop A = Y
Transgenic crop produces an enzyme that adversely affects plant decomposition in soil	Residues of transgenic crop A decompose at the same rate as those of its progenitor B	Enzyme activity will be X units per gram of soil

Risk hypothesis attempt to clarify whether a specified harmful outcome is likely, whereas research hypotheses seek precise quantitative predictions. NOAEC: no observable adverse effect concentration—the highest concentration of a substance that has no observable adverse effect on the species in question.

species. Each stage in the scenario allowed the formulation of a testable risk hypothesis (**Table 2**). Testing of these hypotheses led to the conclusion that introgression of transgenes conferring TuMV resistance into populations of *B. oleracea*, *B. nigra* and *B. rapa* poses low environmental risk in the United Kingdom.[24]

In terms of basic research, risk hypotheses may be uninteresting; however, this does not matter, because interest in risk hypotheses comes from their ability to assist decision-making. Risk hypotheses should be formulated such that their corroboration or falsification provides a good basis for inferring risk from a given scenario. It follows that a risk hypothesis is ineffective if it is not obvious what action to take should the hypothesis be corroborated or falsified. Exposing clear choices is the main function of a risk hypothesis and should not be confused with the "bold anticipations" of hypotheses for basic research. Further examples of risk hypotheses and contrasting basic research hypotheses are given in **Table 3**.

Testing hypotheses. Confidence in the predictions of a hypothesis comes from the rigour with which it has been tested: corroboration by rigorous tests provides more confidence than corroboration by weak tests. Seeking rigorous tests of hypotheses applies to basic research and to risk assessment; however, the type of test that provides rigor tends to differ, with field studies often being the more rigorous tests of hypotheses arising from fundamental ecological theory, and laboratory studies the more stringent tests of risk hypotheses.[12,18,26]

As noted above, rigorous tests provide conditions under which flaws in a hypothesis are most likely to be revealed; we do not search for black swans in places where only white swans are likely to be seen. In ecological research, hypotheses tend to predict the presence of particular phenomena given certain conditions. Many species of crucifer contain compounds called glucosinolates. When leaves of these plants are damaged, for example by herbivory, glucosinolates are degraded to various products, including volatile isothiocyanates. There are various theories about the evolution and ecology of glucosinolates,[27] which lead to predictions that the abundance of herbivores which specialize on crucifers will be positively correlated with leaf concentrations of isothiocyanate-producing glucosinolates, whereas the abundance of generalist herbivores is predicted to be negatively correlated with glucosinolate concentration. Correlations between glucosinolate concentration, isothiocyanate production and herbivore behavior are readily demonstrable in the laboratory;[27-30] however, the laboratory provides an unconvincing test of the hypothesis that glucosinolates control herbivore abundance in the field. In the field, many factors may confound the effects detected in the laboratory: isothiocyanates may be dissipated immediately after release, other volatiles may mask the isothiocyanates, and plants with low concentrations of glucosinolates may have other defenses against herbivores, among other things. A controlled field study with blocks of plants containing different glucosinolate concentrations would therefore be more rigorous than laboratory studies for testing the predicted association between glucosinolate concentration and herbivore abundance. A field study in natural communities of plants with varying glucosinolate concentrations would be a more rigorous test still.

In general, laboratory studies "magnify incidental or trivial factors … indeed, laboratory experiments can likely show some effect of any factor by using sufficiently extreme conditions. Laboratory studies are effective in isolating the response to a factor, but the response may not be ecologically relevant."[31] In the glucosinolate example, laboratory studies show that many insects have strong tendencies to move towards or away from isothiocyanates; however, in the field, the relative abundance of these insects on individuals of a particular plant species is often not predicted by the concentration of isothiocyanate-producing glucosinolates in those plants.[32]

In risk assessment, one is usually testing a hypothesis that a particular phenomenon is unlikely, which is the equivalent of testing the hypothesis that isothiocyanates have no effect on insect behavior. For hypotheses of this kind, laboratory studies are the most rigorous tests because if the effect is not seen in the laboratory, it is even less likely to be seen in the field: if we find no response of herbivores to isothiocyanates in the laboratory, we may have confidence that those compounds are unlikely to be important in the ecology of plant-herbivore interactions. Similarly, if a transgenic protein has no adverse effect on a particular species in the laboratory at concentrations several times greater than those found in the transgenic crop, it is unlikely to have an adverse effect at the lower concentrations in the transgenic plant, and extensive field studies should not be necessary to further characterize risk for regulatory decision-making.[26] Likewise, if it is not possible to create hybrids between a crop and a wild species in the laboratory, perhaps using male-sterile females and embryo rescue, then it is unlikely that the species will hybridize in field experiments.[24] Again, extensive field work should be unnecessary for regulatory decision-making.

A second important point about hypothesis testing is the use of existing data.[33] In basic ecological research, existing data may provide good tests of new hypotheses; however, convincing corroboration of the hypothesis will usually require new experimental tests as well as re-interpretation of existing data. Often existing data are not in a form that provides the best test of a new hypothesis and the data may have been used in formulation of the hypothesis. In risk assessment, on the other hand, it is often possible and desirable to use only existing data to provide satisfactory corroboration of a risk hypothesis. In the case of a transgenic crop producing a non-pesticidal protein, for example, the risk hypothesis that the protein has no adverse effects on wildlife at concentrations in the crop can be tested using data on mode-of-action, amino acid sequence similarity to known toxins, and the taxonomic distribution of similar proteins.[5] Similarly, no new experiments are required to satisfactorily corroborate the risk hypothesis that there will be no gene flow from transgenic maize to wild species in Europe. To some extent, these are examples where formal testing of a risk hypothesis is a device to organize information to show what may be "obvious" and thereby demonstrate that what could be regarded as assumptions are actually well-corroborated hypotheses. It is unlikely that this situation is common in basic research, because what is obvious is unlikely to be interesting.

In summary, the importance of field studies in basic ecological research does not mean that field studies are essential for testing risk hypotheses that have been corroborated by laboratory studies. Field studies for risk assessment are most effective when a risk hypothesis has been falsified in the laboratory; thus, a laboratory study may falsify the hypothesis that a protein has no adverse effects on a certain insect at the concentration in the transgenic crop in question, and a field study could be used to test a new risk hypothesis that the transgenic crop will have no adverse effect on that insect in the field because field conditions mitigate the effect of the protein.[18,33] Finally, a risk assessment is an interpretation of data, not necessarily a collection of new data, and therefore we should not consider a risk assessment unsatisfactory solely because it presents no data collected specifically for the purpose.

Conclusion

If faced with uncertainty about whether the risks of cultivating a transgenic crop are acceptable, it is tempting to request more research. The differences between basic ecological research and risk assessment suggest that more research may do little to improve decision-making,[5] and may make things worse;[6] therefore before requesting more data, the adequacy of the environmental risk assessment should be evaluated:

(1) Has harm been defined? What constitutes a harmful effect is a subjective judgement based on law or other instruments of policy. While research can help to predict the consequences of cultivating a transgenic crop, it cannot conclude whether those consequences are harmful. A discussion about what would be regarded as a harmful effect may be more fruitful than extra research into consequences.

(2) Have clear risk hypotheses been formulated? Risk hypotheses predict the absence of events necessary for harm to occur, and their corroboration or falsification thereby helps decision-making. Complex hypotheses that make precise predictions should be avoided unless decisions depend upon the exact quantification of an effect. Simple hypotheses that make semi-quantitative predictions, such as the effect of a transgenic crop will be no greater than that of a currently acceptable comparator may be more helpful. A discussion about what constitutes a reasonable threshold value for an indicator of potential harm may be more fruitful than more research to increase the precision of predictions of that indicator.

(3) Have the risk hypotheses been rigorously tested? In many cases, laboratory studies are a more rigorous test of a risk hypothesis than are field studies. If an attribute of the transgenic crop or transgenic protein necessary for harm to occur cannot be detected in the laboratory, it is unlikely that it will be detected in a field study. If uncertainty in the level of risk is still too great after laboratory testing, further laboratory testing may be more fruitful than a lengthy, extensive program of field testing.

(4) Have existing data been fully exploited? Assessing risk is not the same as doing an experiment, and the risk posed by a cultivation of a transgenic crop cannot be judged by the willingness of a developer to produce more data. If existing data adequately assess the risk, further research is simply the equivalent of counting white swans, not searching for black ones, and may increase risk if the introduction of beneficial products is delayed.[34]

If the above points are valid, difficult decisions about whether to allow cultivation of transgenic crops may benefit from better risk assessment rather than from extra research. If more data are required, they should be used to test specific risk hypotheses about the likelihood and seriousness of harm. The research should be clearly relevant to decision-making and should not be a substitute for critical thinking about the problems of defining harm and how one should organize information to characterize risk: those problems will still exist regardless of how much research is commissioned.

None of this means that research into the basic ecology of transgenic crops is not worthwhile; however, it may be more rewarding if it seeks to test fundamental ecological theory rather than justifying itself by its usefulness to regulatory risk assessment. Transgenic plants have been used as a research tool to test hypotheses in many areas of plant science, including pathology,[36] physiology,[37] molecular biology,[38] genetics[39] and epigenetics.[40] Use of transgenic plants to test basic ecological theory seems much less prevalent. This may partly reflect the difficulty of obtaining permits for field studies; however, as pointed out above, many field studies of transgenic crops have been carried out[11,21,41] and so lack of opportunity cannot be the complete explanation. Another possible reason why transgenic plants are not used to test ecological theory is that there is simply no ecological theory to test. Some authors[31,42] have criticized ecology for preferring scholastic conceptual debates and detailed descriptions of natural history to the development and testing of empirical theory. There is probably some truth in this criticism; nevertheless, ecological research on transgenic plants does seem to have missed opportunities to test theory.

Many studies that tested a null hypothesis of no difference between, say, a field of a Bt crop and a similar non-Bt crop actually found differences. Once these differences were found, researchers were exceptionally creative at finding explanations for them.[43-45] These explanations may be satisfactory ad hoc answers to specific

questions, but collectively they are unsatisfying because there is no underlying general theory on which the all explanations are based. Much more interesting would be studies that used theory to develop hypotheses about how, say, the effects of Bt toxins on target pest species, or changes in weed communities in transgenic herbicide-tolerant crops, may affect the abundance of other species or functional groups. It seems a missed opportunity that so many field studies of transgenic crops tested null hypotheses of marginal value to risk assessment and regulatory decision-making, instead of developing and testing ecological theory.

Theory that predicts changes in transgenic crops need not be specific to those crops. New general theories of ecosystem structure based on, for example, energy flow through trophic groups, are being developed.[46] It is difficult to think of anything more fundamental to ecology than theories that predict how ecosystems will change if perturbed, and transgenic crops provide an excellent means to test these theories. Testing basic ecological theory with transgenic crops may ultimately be far more valuable to environmental risk assessment than field trials that simply confirm that proteins that are not toxic at high concentrations in the laboratory are also not toxic at lower concentrations in transgenic crops in the field. Well-tested predictive theory about how crop management perturbs trophic functional relationships will be invaluable for making decisions about how to manage agriculture productively and sustainably.

In conclusion, it is important to recognize the differences between environmental risk assessment and basic ecological research because confusing them can lead to ineffective risk assessment and missed opportunities to advance ecological theory. And if faced with a difficult decision about whether to approve the commercial cultivation of a GM crop, better risk assessment may be more helpful than more research.

References

1. Gray AJ. Ecology and government policies: the GM crop debate. J Appl Ecol 2004; 41:1-10.
2. Sanvido O, Romeis J, Bigler F. Ecological impacts of genetically engineered crops: ten years of field research and commerical cultivation. Adv Biochem Eng Biotechnol 2007; 107:235-78.
3. Brookes G, Barfoot P. Global impact of biotech crops: socio-economic and environmental effects, 1996–2006. Ag Bio Forum 2008; 11:21-38.
4. Kessler C, Economidis I. eds. EC-Sponsored Research on Safety of Genetically Modified Organisms: A Review of Results. Luxembourg: European Commission 2001.
5. Craig W, Tepfer M, Degrassi G, Ripandelli D. An overview of general features of risk assessments of genetically modified crops. Euphytica 2008; 164:853-80.
6. Johnson KL, Raybould AF, Hudson MD, Poppy GM. How does scientific risk assessment fit within the wider risk analysis? Trends Plant Sci 2007; 12:1-5.
7. Wolt J, Keese P, Raybould A, Fitzpatrick JW, Burachik M, Gray A, et al. Problem formulation in the environmental risk assessment for genetically modified plants. Transgenic Res 2009; in press.
8. Wolt J, Peterson RKD. Agricultrual biotechnology and societal decision-making: the role of risk analysis. Ag Bio forum 2000; 3:39-46.
9. Lawton JH. Ecology, politics and policy. J Appl Ecol 2007; 44:465-74.
10. Lubechenco J. Entering the century of the environment: a new social contract for science. Science 1998; 491-7.
11. Squire GR, Brooks DR, Bohan DA, Champion GT, Daniels RE, Haughton AJ, et al. On the rationale and interpretation of the farm scale evaluations of genetically modified herbicide-tolerant crops. Phil Trans R Soc Lond B 2003; 1779-99.
12. Raybould A. Ecological versus ecotoxicological methods for assessing the environmental risks of transgenic crops. Plant Sci 2007; 173:589-602.
13. Flew A. How to Think Straight. New York: Promethius Books 1998.
14. Popper KR. The Logic of Scientific Discovery. London: Hutchinson 1959.
15. Popper KR. Objective Knowledge: An Evolutionary Approach. Oxford: Clarendon Press 1972.
16. Greenland S. Induction versus Popper: substance versus semantics. Int J Epidemiol 1998; 27:543-8.
17. Popper KR. Conjectures and Refutations. London: Routledge 2002.
18. Raybould A. Problem formulation and hypothesis testing for environmental risk assessments of genetically modified crops. Environ Biosafety Res 2006; 5:119-25.
19. Caws P. The structure of discovery. Science 1969; 166:1375-80.
20. Braben DW. Pioneering Research: A Risk Worth Taking. Hoboken NJ 2004.
21. Marvier M, McCreedy C, Regetz J, Kareiva P. A meta-analysis of effects of Bt cotton and maize on nontarget invertebrates. Science 2007; 316:1475-7.
22. Garnier A, Lecomte J. Using a spatial and stage-structured invasion model to assess the spread of feral populations of transgenic oilseed rape. Ecol Model 2006; 194:141-9.
23. Wilkinson MJ, Elliott LJ, Allainguillaume J, Shaw MW, Norris C, Welters R, et al. Hybridization between *Brassica napus* and *B. rapa* on a national scale in the United Kingdom. Science 2003; 302:457-9.
24. Raybould A, Cooper JI. Tiered tests to assess the environmental risk of fitness changes in hybrids and wild relatives: the example of virus resistant *Brassica napus*. Environ Biosafety Res 2005; 4:127-40.
25. Wilkinson MJ, Davenport IJ, Charters YM, Jones AE, Allainguillaume J, Butler HT, et al. A direct regional scale estimate of transgene movement from GM oilseed rape to its wild progenitors. Mol Ecol 2001; 9:983-91.
26. Romeis J, Bartsch D, Bigler F, Candolfi MP, Gielkens MMC, Hartley SE, et al. Assessment of risk of insect-resistant transgenic crops to nontarget arthropods. Nat Biotechnol 2008; 26:203-8.
27. Kliebenstein DJ, Kroymann J, Mitchell-Olds T. The glucosinolate-myrosinase system in an ecological and evolutionary context. Curr Opin Plant Biol 2005; 8:264-71.
28. Warton B, Matthiessen JN, Shackleton MA. Glucosinolate content and isothiocyanate evolution—two measures of biofumigation potential of plants. J Agric Food Chem 2001; 49:5344-50.
29. Blande JD, Pickett JA, Poppy GM. A comparison of seriochemically mediated interactions involving specialist and generalist Brassica-feeding aphids and the brachonid parasitoid *Diaeretiella rapae*. J Chem Ecol 2007; 33:767-79.
30. Newman RM, Hanscom Z, Kerfoot WC. The watercress glucosinolate-myrosinase system: a feeding deterrent to caddisflies, snails and amphipods. Oecologia 1992; 92:1-7.
31. Peters RH. A Critique for Ecology. Cambridge: Cambridge University Press 1991.
32. Moyes CL, Collin HA, Britton G, Raybould AF. Glucosinolates and differential herbivory in wild populations of *Brassica oleracea*. J Chem Ecol 2000; 26:2625-41.
33. Romeis J, Lawo NC, Raybould A. Making effective use of existing data for case-by-case risk assessments of genetically engineered crops. J Appl Entomol 2009; 133:571-83.
34. Maund SJ, Sherratt TN, Stickland T, Biggs J, Williams P, Shillabeer N, et al. Ecological considerations in pesticide risk assessment for aquatic ecosystems. Pestic Sci 1997; 49:185-90.
35. Cross FB. Paradoxical perils of the precautionary principle. Wash L Law Rev 1996; 53:851-925.
36. Goodin MM, Zaitlin D, Naidu R, Lommel SA. *Nicotiana benthamiana*: its history and future as a model for plant pathogen interactions. Mol Plant-Microbe Interact 2008; 21:1015-26.
37. Sakamoto A, Murata N. The role of glycine betaine in the protection of plants from stress: clues from transgenic plants. Plant Cell Environ 2002; 25:163-71.
38. Jones-Rhoades MW, Bartel DP, Bartel B. MicroRNAs and their regulatory roles in plants. Annu Rev Plant Biol 2006; 57:19-53.
39. Busov VB, Brunner AM, Meilan R, Filichkin S, Ganio L, Gandhi S, et al. Genetic transformation: a powerful tool for dissection of adaptive traits in trees. New Phytol 2005; 167:9-18.
40. Henderson IR, Jacobsen SE. Epigenetic inheritance in plants. Nature 2007; 447:418-24.
41. Romeis J, Meissle M, Bigler F. Transgenic crops expressing *Bacillus thuringiensis* toxins and biological control. Nat Biotechnol 2006; 24:63-71.
42. Lawton JH. Community Ecology in a Changing World. Oldendorf: Ecology Institute 2000.
43. Candolfi MP, Brown K, Grimm C, Reber B, Schmidli H. A faunistic approach to assess potential side-effects of genetically modified *Bt*-corn on non-target arthropods under field conditions. Biocontrol Sci Technol 2004; 14:129-70.
44. Dively GP. The impact of transgenic VIP3A x Cry1Ab Lepidopteran-resistant field corn on the nontarget arthropod community. Environ Entomol 2005; 34:1267-91.
45. Naranjo SE. Long-term assessment of the effects of transgenic *Bt* cotton on the abundance of nontarget arthropod natural enemies. Environ Entomol 2005; 34:1193-210.
46. Caron-Lormier G, Bohan DA, Hawes C, Raybould A, Haughton AJ, Humphry RW. How might we model an ecosystem? Ecol Model 2009; 220:1935-49.