Economic Profitability Versus Ecological Entropy

Citation

Weitzman, Martin L. 2000. Economic profitability versus ecological entropy. Quarterly Journal of Economics 115(1): 237-263.

Published Version

<http://dx.doi.org/10.1162/003355300554728>

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:3693703>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

ECONOMIC PROFITABILITY VERSUS ECOLOGICAL ENTROPY*

MARTIN L. WEITZMAN

There is a long-standing trade-off in bioculture between concentrating on high-yield varieties and maintaining sufficient diversity to lower the risks of catastrophic infection. The paper uses a simple ecology-based model of endogenous disease to indicate how a local decision to plant more of a widely grown crop creates negative externalities by increasing the probability that new pathogens will evolve to attack the crop globally. Society's basic issue concerns where to locate on an efficiency frontier between economic profitability and a standard formula for ecological entropy—proved here to be a rigorous measure of “generalized resistance” to crop-ecosystem failure.

INTRODUCTION

Privately, the most profitable human use of biological mass is to concentrate intensively on the cultivation of a relatively few high-yield crop varieties. But such widespread monoculture patterns create large target hosts, effectively inviting new potentially lethal pathogens to evolve.

The essence of this trade-off can be understood quite intuitively. Whether in agriculture, aquaculture, horticulture, silviculture, viniculture, dairy-culture, or any other biological culture, high yields come from specialization to cultivating just a few genetically uniform varieties of only a relatively small number of domesticated species. But parasites—in the form of bacteria, viruses, fungi, yeasts, protozoans, nematodes, insects, or others—also come in species and strains. And these parasitic species and strains themselves tend naturally to specialize, in this case to specific hosts. The incredible specificity of most parasite-host relationships is a well-known theme in biology. Other things being equal (and speaking loosely), the more prevalent is the host, the bigger is the size of the evolutionary dining-room area within which the host-specific parasites have leeway to play with new genetic combinations, or to experiment with the increased comparative advantage that comes from specializing to finer-grained subniches within the host organism. An eventual proliferation of

* For useful comments on an earlier version of the paper, I thank Ita Falk, James Hammitt, and David Simpson. One of the anonymous referees provided many detailed suggestions, which, I hope, helped to clarify the present version. The research was financed by a grant from the National Science Foundation.

new parasitic species and strains, including pathogens, is the ages-old natural outcome of selective pressures when the host population is artificially increased.

In this analogy the total biomass of a genetically uniform host is like an island in a sea of other biomass, and a ubiquitous finding of ecology is that, other things being equal, the bigger the size of an island the more species will be located there. Ecologists may debate the ultimate theoretical origins of this rule, but there is little question about its empirical validity. A positive species-area relationship is arguably the most pervasive and the most robust single stylized fact in all of biogeography. Every basic ecology textbook prominently displays a number of scatter diagrams showing that the log of species count is approximately linear in the log of area—over a wide range of contexts and with many different definitions of “area” including, for parasitic species, the total (‘area’) occupied by the host. In the context of this paper it is relevant to note that human-created cropping patterns also constitute an ecosystem, even though the ecological interactions that will be modeled here are restricted to having the same crop, in different locations, being preyed upon by the same mutually shared class of host-specific pathogen predators. Seen this way, the paper is just pushing through to its logical conclusion an ecological interpretation of the history of worldwide bioculture—as representing a very large series of artificial experiments in island biogeography, and therefore as being subject to the same basic underlying rules and regularities as any other experiments in ecology.

When the size of a host population is artificially changed, the number of host-specific parasite species or strains may perhaps remain quasi-fixed for a while at the previous equilibrium level. Thus, in the short run there may appear to be no “disease penalty” for increasing the biomass target area of a cultivated crop. But this is a temporary illusion. Given enough time, forces of immigration and pressures of natural selection will operate among parasite strains, much as Malthusian forces operate within strains. In the long run the equilibrium number of different pathogens cannot be taken as given, but must instead be viewed as endogenously determined by the size of the host population itself.

When humans artificially create or maintain genetically homogeneous host-crop target areas, they are also creating or maintaining breeding grounds with higher probabilities that

potentially lethal pathogens will emerge, some of which, by the laws of chance, could wipe out the very monocultures that spawned or supported them. Thus, there is an inescapable long-run trade-off between the gains of increased crop specialization and the increased risks of lethal infection. Furthermore, this trade-off involves a “disease externality” that makes privately optimal specialization patterns not be socially optimal. When a farmer is deciding which crops to grow, he takes no account of the effects of his local decisions on the global probabilities of crop-ecosystem failures.

Decentralized decision-making farmers will take account of *some* aspects of disease risks of crop failure on other plots—as well as on their own. For example, an individual farmer will allow for the fact that if one crop is wiped out by disease, then prices will be higher for other crops. The problem here is not risk per se, nor being unaware of what other farmers are doing, but the particular negative-externality risk created by the individual farmer’s encouraging the development of more crop-specific parasites globally whenever he decides locally to plant more of that crop. This induced-disease risk represents a genuine externality because there is no automatically self-correcting feedback mechanism, since the individual farmer bears but a negligible fraction of the total disease burden that he may be putting on others every time he makes a cropping decision.

In this particular respect, farmers on their own restricted-access private property behave more like fishermen with free access to common property. For both situations the individual calculates only the effect of his own decision on his own (expected) yield, taking no account of the effects of his decision on the (expected) yields of *other* fishermen or farmers. In what follows throughout the paper, when I talk about “risk” or “externality” it is *this* “disease-externality risk” I will have in mind. To clarify the terminology and to focus the paper very sharply, *there are no other externalities in the model and all of the other kinds of risks are embedded into the problem by just dealing with their expected values.*

The risk of catastrophic crop failure can be lowered by undertaking various preventative or reactive measures, which, of course, is what is done in practice all the time. But the risk can almost never be totally eliminated by some scientific magic bullet, real or imagined. In this spirit, and in such a complicated actuality, I think an appropriate modeling strategy is to take the

probability of failure as being extremely small, but not zero. That is to say, a first-pass model aimed at the big picture might be allowed, by a willing suspension of disbelief, to concentrate on a reduced-form approximation where the probability of any one unit succumbing is parametrically given by some positive epsilon. After the relevant scientific measures have been undertaken, this epsilon approaches zero, but yet never quite gets there. The goal of the paper is to identify general principles of optimal behavior that are robust when epsilon is a small but unknown positive number. While the model treats such an epsilon as if exogenously given, I believe that the basic insights would survive a more complete treatment that might take into explicit account the costs of reducing it.

Note that the cataclysmic collapse of a monoculture ecosystem is not automatically a negligible event just because its probability of occurrence may be an extremely small positive number—since the *consequences* of this occurrence, which presumably include widespread prolonged famine, must be envisaged as an extremely *large* negative utility. Thus, the expected worst-case welfare loss from worldwide crop specialization may actually represent a very significant global externality, even though only very small probabilities are involved.

My point of departure is the idea that humans may be so stuck in the trenches, battling pathogen foes every day, that we fail to stand back a sufficient distance to take full measure of the fundamental character of this long-running war against nature, in which we have been engaged, with accelerating intensity, since neolithic times. Given enough expenditures of money and time, we can usually defeat any one pathogen insurrection, or at least make the probability of crop failure an acceptably small positive number. But is this strategy of putting out fires as they arise a good use of resources overall? Would it not be more prudent to invest, so to speak, in less-flammable biomass proportions right from the beginning? After all, in one way or another, reduced probabilities of extinctions translate, ultimately, into cost savings somewhere in the system.

The purpose of this paper is to analyze the properties of a pervasive long-run relation between social externality risk and private economic return, which arises naturally from applying first principles of probability theory to the simplest possible model of endogenously determined disease and extinction in genetically uniform host populations. I will show, at a high level of abstrac-

tion to be sure, that it is possible to give a concise conceptual interpretation “as if” in terms of a trade-off between net marginal economic profitability, representing the private return, and ecological entropy, standing as the relevant proxy for the social risk externality.¹ This very same net marginal economic profitability can also be given a meaningful interpretation as representing the socially ideal tax or subsidy on a crop, which corrects for the negative or positive marginal contribution of that crop to overall biodiversity.

It then turns out that these corrective taxes or marginal economic profitabilities are related to the socially efficient biomass proportions of the crop by a very simple mathematical formula. This simple rule implies directly a strong rigorously based sense in which, from a worldwide social perspective, it is extremely desirable to maintain as insurance small but viable samples of genetically diversified alternative crops, even though they may be far less profitable than the monoculture-type cash crops currently favored for their profitability. Some further implications and applications of this paper’s framework are also discussed briefly.

ENTROPY AS GENERALIZED RESISTANCE TO ECOSYSTEM FAILURE

In this section I attempt to give a short explanation of the main contribution of the paper to ecology per se. While this brief description might be seen as somewhat of an aside for a “pure” economist reader, my own view is that economic and ecological understandings of the meaning and use of entropy measures complement each other considerably.

Suppose that a given ecosystem community consists of n species. Let the relative biomass² of species i ($i = 1, 2, \dots, n$) be

1. Remember, in the paper the disease-externality risk is simultaneously the only externality and also the only source of uncertainty. Perhaps, therefore, it is slightly more accurate to identify “net marginal economic profitability” as including both the private and the social return (for “ordinary” externalities), but excluding the extra cost associated with catastrophic crop failure, which is captured by the (extra-ordinary) “social risk externality.”

2. “Biomass” is a widely used measure of the total (usually dry) weight of organic material associated with a crop. “Live biomass” refers to the part of the organic material that is currently alive, like the inner bark of a tree, as opposed to the “dead biomass” of the heartwood (or the living mussels, as opposed to the dead mussel shells composing the reef). Of these two biomass measures, we use whichever one is the more economically relevant in a given context.

denoted q_i where $q_i \geq 0$ and

$$(1) \quad \sum_{i=1}^n q_i = 1.$$

Perhaps the most widely used measure of community-level biodiversity is the Boltzmann-Shannon entropy formula,

$$(2) \quad H' = - \sum_{i=1}^n q_i \log q_i.$$

The above formula can be defended on several plausible grounds. Sometimes a series of informal arguments are made “by analogy with” thermodynamics or information theory.³ More typically, it is shown heuristically that actually using H' as a criterion gives intuitively desirable results in a number of hypothetical and actual examples. However, the only rigorous justification advanced in the ecology literature is that (2) represents the unique functional form allowing consistent aggregation over classification levels. While this is truly a desirable and an important property for a biodiversity measure to possess, it is perhaps even more important to be aware that the “ecological entropy” formula as it currently stands has no *direct* ecological interpretation.⁴

In this respect, economics is somewhat like ecology. The entropy formula (2) is used in economics primarily for its congenial analytical properties as a measure of inequality.⁵ But there is no *direct* interpretation of H' in terms of some underlying economic process.

Coming at this issue of directly interpreting H' from another direction, conventional ecological wisdom traditionally holds that more diverse communities are in some sense more stable. Yet, the

3. In statistical mechanics the probabilities refer to the relative numbers of microstates in position-momentum phase space that lie behind observable macrostates, like temperature or pressure, and as such entropy is widely interpreted as measuring the overall “disorderliness” of a thermodynamic system; see, e.g., Sears [1953]. In information theory, the entropy formula is widely used to quantify the information content of a message source, as measured by the minimum required channel capacity for its meaningful transmission; see, e.g., Raisbeck [1964].

4. On this subject, Pielou [1977], who is a widely recognized authority, writes: “There has been much debate on whether H' is a suitable measure of ecological diversity. The fact that it measures “information” and “entropy” is beside the point; these fashionable words have been bandied about out of their proper context (the mathematical theory of information) and have led to false analogies that produced no noticeable advance in ecological understanding.”

5. See, e.g., Cowell [1995] or Theil [1967].

formal connection has remained elusive.⁶ The basic contribution of the present paper to ecology proper is a demonstration that there exists an internally consistent ecology model showing that the ecological entropy formula (2) can, at least in principle, be interpreted rigorously as an exact theoretical measure of a particular type of ecosystem stability—essentially the probability that the system can successfully resist a catastrophic “extinction shock” from endogenously generated independently acting pathogens. Furthermore, it will be shown that there is an exact sense in which entropy is “partitionable” to any level of inclusivity of the system, and hence might be considered a “generalized measure” of resistance to extinction failure of this kind.

The finding that entropy is interpretable as an index of a certain kind of ecosystem robustness may be of some interest worth noting for ecology, since the concept has previously lacked any direct interpretation. However, my main purpose in the paper is not so much to demonstrate this interpretation as to *use* it in characterizing the kind of macro-level long-run aggregate trade-off that the world faces in choosing between high crop yields from specialization, on the one side, and, from the other side, the associated vulnerability to failure of an artificially selected monoculture-type ecosystem.

COMPETITIVE EQUILIBRIUM AND EXTERNALITY-CORRECTING TAXES

The primary aim of this paper is to explore the nature and significance of a particular form of externality—called here a “disease externality.” Before proceeding further, it would be well to lay out a general framework now, the better to understand later where this particular type of externality fits into the standard analysis.

Throughout the remainder of this paper, the word “crop” is intended to be used in a generic sense, to stand for some artificially cultivated variety (“*cultivar*”) of domesticated plant, animal, yeast, or bacteria, which is, or can be, raised in agriculture, aquaculture, horticulture, silvaculture, vinoculture, dairy-culture, or any other form of bioculture.

We begin at a global level of abstraction. Suppose that the total number of potential domesticates on the planet, which are or

6. For a textbook discussion of the issues and debates, see, e.g., Chapter 23 of Begon, Harper, and Townsend [1990] or Chapter 17 of Pianka [1994].

might be considered potentially useful, is denoted by n , which is likely to be a large number absolutely, yet is probably quite small by comparison with the total number of all wild types of species and varieties existing in the natural world.

Let the worldwide annual *biomass harvest* of crop i ($i = 1, 2, \dots, n$) be denoted B_i , with associated n -vector $\mathbf{B} = (B_i)$. Suppose that the *utility* of biomass harvest $\mathbf{B} \geq 0$ is given by the smoothly differentiable increasing concave function,

$$(3) \quad U(\mathbf{B}).$$

Let the *price* of a unit of crop i ($i = 1, 2, \dots, n$) be denoted π_i , with associated n -vector $\pi = (\pi_i)$. Denote the *demand function* for $\mathbf{B} \geq 0$ as

$$(4) \quad \mathbf{D}(\pi),$$

where, for all values of $\mathbf{B} \geq 0$ it holds that

$$(5) \quad \mathbf{D}(\mathbf{U}'(\mathbf{B})) = \mathbf{B}.$$

It will be convenient to treat externalities here as a difference between social and private costs. Let the *private cost function*, which represents the least-cost way of producing the biomass vector $\mathbf{B} \geq 0$ in a private, decentralized, competitive economy, *taking no account of externalities*, be given by the smoothly differentiable increasing convex function,

$$(6) \quad \mathbf{C}(\mathbf{B}).$$

Let the corresponding *social* cost function, which represents the efficient least-cost ideally planned way of producing the biomass vector $\mathbf{B}$, as if *internalizing all externalities*, be denoted

$$(7) \quad \Psi(\mathbf{B}).$$

The decentralized *private supply function* is

$$(8) \quad \mathbf{S}(\pi),$$

where for all values of $\mathbf{B} \geq 0$, it must hold that

$$(9) \quad \mathbf{S}(\mathbf{C}'(\mathbf{B})) = \mathbf{B}.$$

The *decentralized private equilibrium* here occurs at prices $\underline{\pi}$ and quantities $\underline{\mathbf{B}} \geq 0$, where

$$(10) \quad \mathbf{D}(\underline{\pi}) = \mathbf{S}(\underline{\pi}) = \underline{\mathbf{B}}.$$

What might be called the *private net economic value* of biomass harvest $\mathbf{B} \geq \mathbf{0}$ is

$$(11) \quad V(\mathbf{B}) \equiv U(\mathbf{B}) - C(\mathbf{B}),$$

while the *social value* of biomass harvest $\mathbf{B} \geq \mathbf{0}$ is

$$(12) \quad V^*(\mathbf{B}) \equiv U(\mathbf{B}) - \Psi(\mathbf{B}).$$

The decentralized private-equilibrium biomass vector $\mathbf{B} = \mathbf{D}(\pi) = \mathbf{S}(\pi)$ can be interpreted *as if* maximizing the private net economic value expression (11), which results in the standard duality conditions,

$$(13) \quad \mathbf{V}'(\mathbf{B}) \leq \mathbf{0}, \quad \mathbf{B} \geq \mathbf{0}, \quad \mathbf{V}'(\mathbf{B}) \cdot \mathbf{B} = 0.$$

However, in the presence of externalities $\mathbf{B}$ is *not* socially optimal. The *socially optimal* value of $\mathbf{B}$ would instead maximize the *social* value expression (12).

It is convenient here to rewrite (12) as

$$(14) \quad V^*(\mathbf{B}) \equiv V(\mathbf{B}) + E(\mathbf{B}),$$

where

$$(15) \quad E(\mathbf{B}) \equiv C(\mathbf{B}) - \Psi(\mathbf{B})$$

represents the *social-externality value* of $\mathbf{B}$. The purpose of using the form (14) is to decompose $V^*(\mathbf{B})$ into two conceptually distinct components. The first component $V(\mathbf{B})$ represents the *private-economic value* of $\mathbf{B}$. The second component $E(\mathbf{B})$ represents the *social-externality value* of $\mathbf{B}$.

The *socially optimal* biomass vector $\mathbf{B}^*$ maximizes (14), resulting in the first-order condition,

$$(16) \quad \mathbf{V}'(\mathbf{B}^*) = \mathbf{E}'(\mathbf{B}^*).$$

An important role in the paper will be played by the vector,

$$(17) \quad \tau^* \equiv \mathbf{V}'(\mathbf{B}^*) (= -\mathbf{E}'(\mathbf{B}^*)).$$

The vector τ^* is interpretable from (17) and (11) as being the set of socially optimal *net marginal economic profitability* coefficients. An important equivalent interpretation is that τ^* represents the set of socially optimal *externality-correcting taxes*—in the sense that imposing taxes τ^* would indirectly induce the socially optimal equilibrium response,

$$(18) \quad \mathbf{B}^* = \mathbf{D}(\pi^*) = \mathbf{S}(\pi^* - \tau^*),$$

for the socially optimal equilibrium price vector,

$$(19) \quad \pi^* \equiv \mathbf{D}^{-1}(\mathbf{B}^*).$$

The coefficient τ_i^* in this paper represents the ideal “monoculture tax” (if positive) or “diversity subsidy” (if negative), which corrects for the marginal disease-externality effects of growing one more unit of crop i . Equivalently, τ_i^* is simply the (socially ideal) net marginal economic profitability of raising one more unit of crop i . A negative value of τ_i^* corresponds here to a currently unprofitable crop, which may yet be worthwhile growing for its value as a backup alternative if some mainline cash crops were to fail.

In this paper the only externality being considered is the “disease-externality risk” of crop infection. What stands in for E in expressions (14) and (16) will then essentially turn out to be ecological entropy, which, it will later be proved, is here the appropriate externality measure of generalized resistance to crop-ecosystem failure. For cost-benefit or decentralization purposes *on the margin*, τ will here represent the relevant *private-economic* components of welfare for small changes, while corresponding changes in E will represent the relevant *social-externality* components of welfare.

A MODEL OF ENDOGENOUS DISEASE

The spirit of this model is in the tradition of a macro-level style of analysis that tolerates assumptions of extreme symmetry on the micro level—the better to focus, at the appropriate large scale for viewing the big picture, on the overarching relationship between important aggregates. The treatment here of endogenously determined disease and extinction is not at all fancy, being just the simplest imaginable application of basic probability laws. In a sense, the underlying notion of a distinctive crop “variety” is allowed to be sufficiently elastic to *permit* the sharp symmetry and independence assumptions of the paper to hold as approximations. The model is so highly aggregated that specifications of timing and dynamics are necessarily oversimplified. At the end of the day, my only excuse for such a formulation is that it may deliver a conceptually useful benchmark that might otherwise be obscured by having too many local details. Without further apologies for simplicity or abstraction, I proceed directly to the model.

In what follows, we assume that the “unit of effective biomass” has been standardized across all crops so that the carrying capacity for parasites is symmetrically identical. With respect to location, a crop is presumed to be randomly dispersed geographically in a noncontiguous checkerboard-like micro-pattern, so that, from the lofty vantage of the macrolevel, the grid squares look small enough to allow abstraction away from the issue of “contact contagion” arising because adjacent plots grow identical crops. For convenience, in what follows, suppose that every grid-square pure stand of a crop is of unit-biomass size.

In the short run, let the total biomass of crop i be fixed at

$$(20) \quad \bar{B}_i$$

while the number of potentially lethal strains of i -relevant pathogens is fixed at

$$(21) \quad \bar{S}_i$$

The model of disease here is so crudely basic that it abstracts away from most dynamic issues, such as the time trajectories of contagion, infection, spread, and so forth. Essentially, every member of the host population in this model is postulated to have symmetrically identical exposure and resistance. Pathogens are endogenously generated with a delay lag of one period. All of the subsequent action then plays itself out as if it were occurring instantaneously.

The simplest way of modeling crop mortality is to assume that “resistance” is random, or, what is essentially the same thing in this model, that pathogen-induced deaths are independently distributed across biomass grid squares. Let the probability of *one* particular i -relevant pathogen becoming virulent and wiping out completely *one* particular stand of a unit of biomass of crop i be

$$(22) \quad \epsilon.$$

Then the corresponding probability that $\bar{S}_i$ independently acting pathogens will destroy completely one particular stand of a unit of biomass of crop i is

$$(23) \quad 1 - (1 - \epsilon)^{\bar{S}_i}.$$

Finally, the probability of a catastrophic worldwide extinction of crop i is

$$(24) \quad [1 - (1 - \epsilon)^{\bar{S}_i}]^{\bar{B}_i}.$$

Note that as $\bar{S}_i$ is made larger, or $\bar{B}_i$ is made smaller, expression (24) increases monotonically toward one, whereas $\bar{B}_i$ made larger, or $\bar{S}_i$ made smaller, causes (24) to decline monotonically toward zero. It is not immediately obvious what would happen if both $\bar{S}_i$ and $\bar{B}_i$ were simultaneously increased or decreased. Presumably, the outcome would depend on the nature of the underlying relationship between them, to which topic we now turn.

Let us envision the model dynamics as if unfolding in three discrete periods. Thus far, in the *short run*, which here stands for the past, B_i has been exogenously fixed at $\bar{B}_i$, and S_i has been exogenously fixed at $\bar{S}_i$. Suppose, next, that in the *medium run*, symbolizing the present, B_i is allowed to vary artificially by human intervention, while S_i remains fixed at $\bar{S}_i$. In the *long run*, standing for the future, there is postulated to be a natural parasite-host reactive relationship of the form,

$$(25) \quad S_i = k(B_i)^z,$$

where k and z are positive constants. (I think that such a naively discrete timing sequence captures fairly the analytical essence of the problem, but it does abstract away from potentially complicated dynamic interactions and strategies.)

There is an enormous amount of ecology literature justifying, as a reduced-form empirical approximation holding over a wide range of contexts, a log-log linear relationship of exactly the form (25) between species and area.⁷ In such studies, “area” may have a very general meaning, including, for parasites, how widespread is the host; i.e., in how many grid squares on a map is the host found. The model here is implicitly dealing with a long run of sufficient duration that forces of immigration and pressures of natural selection on micro-parasites, including pathogens, permit (25) to operate as a valid approximation. While it was never contemplated by ecologists that (25) might be applied on the scale envisioned by this paper, I am here making that leap of faith.

We now have a crude but workable theory of endogenous disease that is at least not openly inconsistent with the highly stylized ecological facts summarized by (25). Combining (25) with (24), the long-run *endogenously determined* probability of the

7. See, for example, Begon, Harper, and Townsend; [1990], Conner and McCoy [1979]; MacArthur and Wilson [1967]; Meffe and Carroll [1994]; Pianka [1995]; and Huston [1994]. Values of z for islands typically are within a relatively narrow range from about 0.24 to about 0.34.

complete extinction of crop i is

$$(26) \quad P_i(B_i) = [1 - (1 - \epsilon)^{k(B_i)^z}]^{B_i}.$$

Expression (26) is a measure of the long-run *vulnerability* of crop i . Forces of natural selection favor the evolution of parasites that do not kill off their hosts too rapidly or too completely. In this spirit, (26) describes the probability of pathogens randomly or “inadvertently” becoming sufficiently virulent that they destroy their host species and, with the host, themselves.

From applying l'Hôpital's rule multiple times to the logarithm of the right-hand side of (26), we obtain the following two simple but important inferences about long-run limiting crop vulnerability at the two biomass extremes:

$$(27) \quad \lim_{B \rightarrow 0} [1 - (1 - \epsilon)^{k(B)^z}]^B = 1,$$

and

$$(28) \quad \lim_{B \rightarrow \infty} [1 - (1 - \epsilon)^{k(B)^z}]^B = 1.$$

Condition (27) means that when there are very small numbers in a host population, then there is also a very great vulnerability of losing the entire population to pathogens. Essentially, this occurs because not many members then need be killed to cause extinction. In the “contest of small numbers” between an exogenously determined smaller host population and the endogenously generated smaller number of corresponding pathogens, the pathogens “prevail” in the sense of extinguishing the host. It should be appreciated that such a result is not obvious. When $\bar{B}_i$ alone is made to approach zero, then expression (24) approaches one, the converse of the idea that there is “safety in numbers” for the host. What is *not* so clear, however, is what should happen to expression (24) as *both* $\bar{S}_i$ and $\bar{B}_i$ are simultaneously decreased, since the disease threat ultimately causing extinction is thereby eliminated altogether in the limit. Viewed in this light, result (27) may be seen as extending or strengthening, to a situation that takes explicit account of endogenously induced diseases, our natural biological intuitions about the importance of having some minimum viable population size to avoid random extinctions.

Equation (27) describes an extinction event that may occur in the wild, and in fact happens continually on an evolutionary time

scale. By contrast, condition (28) is more brazenly counterintuitive to ecological principles because it is describing an artificial experiment that has no natural counterpart on nature's time scale of evolutionary interactions—even, possibly, for relatively quick-reacting pathogens such as bacteria or viruses.

The limiting operation in (28) is describing a delayed density-dependent situation where humans in historical time are very rapidly forcing a crop to exist on a biomass scale far beyond the level that nature, who would “react” to a slower process by continuously “creating” new pathogen strains, would ever allow to transpire in the wild. But, it turns out here, nature cannot be pushed indefinitely far in this direction. *Eventually*, the pathogens will catch up, and then with a lagged vengeance, to the artificially expanded host biomass. Equation (28) means that for any positive ϵ , there is some sufficiently large B_i , which will make crop i crash completely after the pathogens react.

Thus, the complete long-run extinction of crop i becomes a certainty in the limit as B_i is made indefinitely large. If all of world agriculture were devoted to growing just one food cultivar, then a delayed catastrophic extinction failure would almost be assured in the long run, as the number of potentially lethal pathogens “locking in” over time to this one host's genetically uniform biomass would become correspondingly large, eventually overwhelming every member of the host population.

There is no safety in numbers here. Just the opposite is true—very high forced numbers will kill off a population in the long run as surely as very low numbers. In the “contest of big numbers” between an exogenously determined larger host population, *each* of whom must be independently killed off for there to be an extinction event here, and the endogenously generated larger numbers of corresponding pathogens, the pathogens will ultimately “prevail” by truly killing off every single last host member. Therefore, depending on the lag parameters, it is theoretically conceivable that humans may inadvertently be constructing some kind of a delayed-fuse time bomb by offering up so much tempting monoculture biomass to viruses, bacteria, and various other mutation-prone reactive pathogens. If such an effect is empirically relevant, I think it is less likely to actually manifest itself in the real world by anything so dramatic as the food supply evaporating some day. Rather, I believe it is more likely to show up as a gradual future increase in real crop-security-related expendi-

tures, somewhat analogous to the steady rise over time of medical spending.

Let us next examine from (26) the behavior of P_i as a (continuous) function of B_i . Implicit in the following exercises, ϵ is taken to be some very small positive number.

From (27) we know that $P_i(0) = 1$. As B_i is increased up from zero, the function $P_i(B_i)$ declines rapidly toward zero. The drop in $P_i(B_i)$ from $P_i(0) = 1$ to “near-zero” for $B_i > 0$ is more precipitous as ϵ is smaller. (In the limit as ϵ is made to approach zero, $P_i(B_i)$ drops instantaneously and discontinuously from $P_i(0) = 1$ to $P_i(B_i) = 0$ for all $B_i > 0$.)

For small positive ϵ , as B_i is increased further, $P_i(B_i)$ hovers slightly above zero for a wide range of $B_i > 0$. The smaller is ϵ , the closer is $P_i(B_i)$ to zero, and the wider is the range over which the function is nearly zero. Then, quite abruptly, at some point as B_i is increased further, the function (26) relatively suddenly climbs rapidly toward one and thereafter remains very near to but just below one, approaching ever closer asymptotically. The transition of $P_i(B_i)$ from “near zero” to “near one” is more precipitous as ϵ is smaller. In the limit as ϵ is made to approach zero, the derivative of (26) at the “transition point” approaches infinity.

I think there are two basic messages that emerge from examining carefully the properties of the function $P_i(B_i)$ defined by (26) for small positive ϵ . The relatively sudden transition from $P_i(0) = 1$ to $P_i(B_i) \sim 0$ for $B_i > 0$ is sufficiently dramatic for small ϵ that the model seems to be warning us to keep away from the $B_i = 0$ region. At the opposite extreme of large and increasing B_i , the sudden transition from the $P_i \sim 0$ phase to the $P_i \sim 1$ phase seems also to be trying to convey an intended message. I think the model here may be trying to alert us that, in a real world having delayed-response lags for the evolutionary development of new pathogens, monoculture-like patterns of extreme crop specialization in the recent past and in the present might conceivably be building surreptitiously toward some nasty surprises, which may only become apparent in the future.

Taken together, equations (26), (27), and (28) presage the main theme of the paper. Pathogens “prevail” in driving their hosts (and themselves) to long-run extinction at the extremes of host populations, while crops and their patrons “prevail” in the middle. Furthermore, the middle here may quite possibly be comfortably large for small ϵ , but then, to the extent that this is true, it is also true that the transition to extinction is uncomfort-

ably abrupt. The overall message, I think, is that as patrons it greatly behooves us to have sufficient *balance* built into the overall design of our cropping ecosystems, so that the impacts of long-run extinction events corresponding to (27) or (28) are somehow marginalized. In a sense, the rest of the paper builds upon this basic theme, the core elements of which flow so naturally from simply noting the consequences of (27) and (28), along with examining the behavior of (26) for small ϵ .

From (26) it follows that the long-run endogenously determined probability of an ultra-catastrophic lagged mass extinction of *all* crops is

$$(29) \quad P(\mathbf{B}) = \prod_{i=1}^n P_i(B_i).$$

Expression (29) is a measure of the overall long-run vulnerability of the planetary cropping ecosystem to catastrophic extinction failure. This kind of a cataclysmic worldwide disappearance of all domesticated varieties represents a biological extinction event that, by standards of natural history, is actually relatively small in terms of numbers of species or races involved, but would be about as horrific for humanity as a medium-large-sized asteroid hitting the earth. A catastrophic collapse of the cropping ecosystem is a possibility that humankind does not want to begin to contemplate, even remotely. Thus, any rational person would insist that world bioculture be designed with a *very* large safety margin, so that expression (29) is kept at a *very* small value.

I believe that the concept of the vulnerability of an ecosystem to lagged "extinction failure" may serve as a useful guide for thinking about cropping-biomass design issues for much less extreme situations than total worldwide collapse. While the model has thus far been phrased in terms of a global extinction of all crops, it will be shown later that exactly the same methodology pertains to *any subset* of crops, whatever the level of application. Because of this important corollary to the theory, it will be possible to interpret (29) as a *generalized measure of the probability of failure*, applicable equally for analyzing any combination of extinction events. Keeping expression (16) low may then be seen as a proxy for maintaining an ecosystem in sufficient balance, at *all* levels, so that damaging extinctions are *generally* unlikely, or at least are less likely in general than they might otherwise be.

This paper does not explicitly model the externality costs of

extinction failures. The formulation here does not lay out the alternatives in a standard risk-versus-return framework based on expected utility theory. There are two reasons for this. First of all, the standard expected-utility formulation is analytically intractable here, resulting in an uninterpretable mess, when what we are really after is a few robust principles to help guide us through the morass. Second, and more substantively, it is not the least bit clear that the underlying axioms justifying the expected-utility hypothesis should apply here at all.⁸

The kind of “risk” that this paper is concerned with is the low-probability, high-impact, nonstandard possibility of some monoculture-type cultivars being so decimated by lagged endogenously generated diseases that, for practical purposes of further usage, they are rendered extinct. For such an extreme event, it is unclear whether the expected-utility approach has more claim to primacy than any other intuitively plausible framework. In the paper this very complicated set of issues is approached indirectly. The approach taken here is simply to examine on the margin the trade-off between a generalized measure of economic welfare and a generalized measure of ecosystem vulnerability. In such a context, the full justification for viewing expression (29) as a generalized measure of vulnerability to extinction failure must await Corollary 2, which, unfortunately, cannot meaningfully be pulled out of its logical sequence and developed independently here.

At this point in the paper, the most immediate task is to specify those basic underlying principles of efficient design that might indicate how high values of (11) may be maintained while ensuring that (29) is tolerably low. For now, the construction of the efficient-possibilities frontier showing how $V(\mathbf{B})$ trades off against $P(\mathbf{B})$ is taken as an intuitively plausible desideratum in its own right, since it can be viewed as representing a useful input to any subsequent decision-making process.

Thus, the approach of this paper is to explain, in the spirit of how a technocratic consultant might present options to a client who will make the final decisions, the nature of the choice between a generalized measure of economic welfare and a generalized measure of resistance to ecosystem failure. As it turns out, the relevant generalized measure of resistance to ecosystem failure, which proxies the unlikelihood of damaging extinctions when ϵ is

8. Some discussion of this set of issues is contained in Schwarz [1998].

small, will be none other than “ecological entropy.” The ecological content of the main theorem then can be seen as indicating an exact limiting sense in which the ecosystem-balance measure H' (formula (2)) may be considered to be a decreasing monotone transformation of ecosystem vulnerability P (formula (29)).

THE BASIC RESULT

Consider a long-run efficient-biomass trade-off problem of the canonical form, *mimimize* over $B_i \geq 0$

$$(30) \quad \prod_{i=1}^n [1 - (1 - \epsilon)^{k(B_i)^z}]^{B_i},$$

subject to

$$(31) \quad V([B_i]) \geq \bar{V},$$

where the parametrically fixed $\bar{V}$ represents some predetermined level of economic value.

It is not difficult to show that the solution of problem (30)–(31) can be written as a set of well-behaved parametric functions of the form,

$$(32) \quad \tilde{B}_i(\bar{V}; \epsilon),$$

for all $i = 1, 2, \dots, n$.

Next, as seems appropriate to the context of a problem where the failure rate is being taken as an extremely small but positive parameter, let ϵ approach zero.⁹ Define

$$(33) \quad B_i^*(\bar{V}) \equiv \lim_{\epsilon \rightarrow 0} \tilde{B}_i(\bar{V}; \epsilon).$$

Finally, define $\{B_i^*\}$ to be an *efficient* biomass allocation (at very low individual failure rates) if it is efficient for some $\bar{V}$; i.e., if there exists a value $\bar{V}$ such that

$$(34) \quad B_i^* = B_i^*(\bar{V})$$

9. Mathematically, what I have in the back of my mind is a corresponding offsetting change in the measurement units of $\mathbf{B}$, which leaves identical the overall probabilities of extinction. A rigorous treatment developed along these more complicated lines gives essentially the same result as the simpler convention adopted in the paper of just letting ϵ alone vary.

for all $i = 1, 2, \dots, n$. Let the corresponding set of marginal profitability coefficients, from (17), be denoted $\{\tau_i^*\}$.

The corresponding efficient *aggregate* biomass $\hat{B}$ is defined as

$$(35) \quad \hat{B} \equiv \sum_{i=1}^n B_i^*$$

while the corresponding set of *efficient biomass proportions* $\{q_i^*\}$ is defined as

$$(36) \quad q_i^* \equiv B_i^*/\hat{B}.$$

The following theorem is the main result of the paper.

THEOREM: Any set of *efficient biomass proportions* $\{q_i^*\}$ must be supported by some positive value of λ as the solution of the following problem: *maximize* over $q_i \geq 0$

$$(37) \quad \sum_{i=1}^n \tau_i^* q_i + \lambda \left[- \sum_{i=1}^n q_i \log q_i \right],$$

subject to

$$(38) \quad \sum_{i=1}^n q_i = 1.$$

Proof of Theorem. Define

$$(39) \quad \tilde{B}(\bar{V}; \epsilon) \equiv \sum \tilde{B}_i(\bar{V}; \epsilon).$$

In what follows, the explicit dependence of $\tilde{B}$ on $\bar{V}$ and ϵ is suppressed for notational convenience.

Consider next the following optimization problem expressed in biomass *proportions*, which, given $\tilde{B}$, is equivalent to (30)–(31): *minimize* over $\{q_i\} \geq 0$

$$(40) \quad \prod_{i=1}^n [1 - (1 - \epsilon)^{k(\tilde{B}q_i)^z}]^{\tilde{B}q_i}$$

subject to

$$(41) \quad V(\{\tilde{B}q_i\}) \geq \bar{V},$$

and

$$(42) \quad \sum q_i = 1.$$

From a Taylor series approximation, we have that

$$(43) \quad 1 - (1 - \epsilon)^{k(\tilde{B}q_i)^z} = \epsilon k(\tilde{B}q_i)^z [1 + O_i(\epsilon; q_i)],$$

where, because the function being approximated is regular in its interior domain, we have that the first-order term $O_i(\epsilon; q_i)$ converges uniformly to zero in the limit as ϵ approaches zero.

Substituting from (43), and making use of (38), the objective function (40) becomes transformed into

$$(44) \quad \epsilon^{\tilde{B}} k^{\tilde{B}} \tilde{B}^{z\tilde{B}} \prod_{i=1}^n [(1 + O_i(\epsilon; q_i)) \cdot (q_i)]^{z\tilde{B}q_i}.$$

Next, dividing the objective function (44) by the positive expression,

$$(45) \quad \epsilon^{\tilde{B}} k^{\tilde{B}} (\tilde{B})^{z\tilde{B}},$$

will not effect the optimization problem because none of the terms in (45) is dependent upon $[q_i]$. Thus, the form (44) may be replaced by an equivalent objective function of the form,

$$(46) \quad \prod_{i=1}^n (1 + O_i(\epsilon; q_i))^{z\tilde{B}q_i} \prod_{i=1}^n (q_i)^{z\tilde{B}q_i}.$$

But note that

$$(47) \quad \lim_{\epsilon \rightarrow 0} \prod_{i=1}^n (1 + O_i(\epsilon; q_i))^{z\tilde{B}q_i} = 1.$$

Therefore, as ϵ is made to approach zero, $\tilde{B} \rightarrow \hat{B}$, and the solution of the original problem (30)–(31) becomes equivalent to the solution of the problem of minimizing

$$(48) \quad \prod_{i=1}^n (q_i)^{z\hat{B}q_i}$$

subject to (41), for $\tilde{B} \equiv \hat{B}$, and to (42).

But minimizing an objective function of the form (48) is equivalent to minimizing the logarithm of (48), which is equivalent to minimizing

$$(49) \quad \sum q_i \log q_i$$

which in its turn is equivalent to *maximizing* the entropy expression (2).

Summarizing a long chain of reasoning, to this point we have demonstrated that the $[q_i^*]$ defined by (36) must maximize (2) subject to (41) for $\tilde{B} = \hat{B}$, and subject to (42). It is readily shown that this problem has a very well-behaved convex structure with an interior solution. Uniqueness of solution is guaranteed by the strict concavity of (2). All that remains is to specify the corresponding dual multipliers.

Let

$$(50) \quad \mu > 0$$

represent the positive shadow price on the limiting version of inequality (41) relative to the objective function (2). Then define

$$(51) \quad \lambda \equiv 1/\mu\hat{B}.$$

Finally, a routine application of duality theory implies that $[q_i^*]$ must be the unique solution of (37)–(38). ■

TWO USEFUL COROLLARIES

Before proceeding to a discussion, it will be helpful first to lay out formally two basic implications of the main theorem.

COROLLARY 1 (CLOSED-FORM CHARACTERIZATION OF EFFICIENT BIOMASS PROPORTIONS).

For all $i = 1, 2, \dots, n$, the solution of problem (37)–(38) is the closed-form expression,

$$(52) \quad q_i^* = \frac{e^{\tau_i^*/\lambda}}{\sum e^{\tau_i^*/\lambda}}.$$

Alternatively, the relationship between any two socially optimal “disease-externality taxes” (or, if negative, “diversity subsidies”) is of the form,

$$(53) \quad \tau_j^* - \tau_i^* = \lambda (\log q_j^* - \log q_i^*).$$

Proof of Corollary 1. If θ is the Lagrange multiplier for equation (38), relative to the objective (37), then the necessary and sufficient first-order conditions for the problem are

$$(54) \quad \tau_i^* = \lambda \log q_i^* + \lambda + \theta.$$

The desired result (52) follows directly from combining the

exponentiated version of condition (54) with equation (38). Condition (53) follows directly from (52). ■

COROLLARY 2 (INVARIANCE OF EFFICIENT BIOMASS PROPORTIONS TO DECOMPOSITION).

There are $(2^n - 1)$ possible subsets that can be formed out of n elements (excluding the null set). Consider *any* such subset economy/ecology consisting of m cultivated varieties, without loss of generality here indexed by $j = 1, 2, \dots, m$, where $m \leq n$. Let $\{p_j^*\}$ be the solution of the following subproblem: maximize over $p_j \geq 0$

$$(55) \quad \sum_{j=1}^m \tau_j^* p_j + \lambda \left[- \sum_{j=1}^m p_j \log p_j \right],$$

subject to

$$(56) \quad \sum_{j=1}^m p_j = 1.$$

Then, for all $j = 1, 2, \dots, m$,

$$(57) \quad p_j^* = \frac{q_j^*}{\sum_{j=1}^m q_j^*}.$$

Proof of Corollary 2. From Corollary 1,

$$(58) \quad p_j^* = \frac{e^{\tau_j^*/\lambda}}{\sum_{j=1}^m e^{\tau_j^*/\lambda}},$$

for $j = 1, 2, \dots, m$, while

$$(59) \quad q_j^* = \frac{e^{\tau_j^*/\lambda}}{\sum_{i=1}^n e^{\tau_i^*/\lambda}},$$

and

$$(60) \quad \sum_{j=1}^m q_j^* = \frac{\sum_{j=1}^m e^{\tau_j^*/\lambda}}{\sum_{i=1}^n e^{\tau_i^*/\lambda}}.$$

Now combine (58) with (59) and (60) to yield the desired expression (57). ■

DISCUSSION

The main theorem expresses the idea that when crop mortality is endogenously determined, efficient combinations of social-externality “risk” and private-economic “return” in this context may be conceptualized as if optimal crop proportions are being generated by maximizing net marginal economic profitability, represented by the coefficients-vector τ^* , plus the parameter λ times ecological entropy H' . (The economic-profitability coefficient τ_i^* is equivalently interpretable as representing the socially ideal “disease-externality tax” levied per unit of crop i . A negative tax here is understood as representing a “diversity subsidy,” which assesses the marginal contribution of crop i to overall biodiversity.) The weight λ reflects the value that society places on overall ecosystem security relative to the more standard type of economic welfare function, measured in dollars, which appears throughout the paper as V (or $\tau^* \equiv \mathbf{V}'(\mathbf{B}^*)$).

Corollary 2 shows that the very same way of thinking about efficient biomass proportions gets replicated at all possible levels of subgroup combinations. For any possible combination of crops here lumped together as a decision unit, a decentralized decision-maker can think in terms of exactly the same kind of balance between the relevant decentralized component of standard economic welfare, represented by the applicable subset of net marginal profitability coefficients, and the relevant decentralized component of crop security, represented by the entropy of crop proportions—with both components pertaining *only to this particular* decision unit.

It is important to realize that, among all possible candidates for a diversity function, just the entropy formula is “partitionable” in this unique manner, which allows efficient biomass proportions to be consistently replicated for every possible subset by using the same diversity function of the proportions of that subset alone. Of course, this partitioning result depends on the benchmark assumption of independent probabilities, but it does indicate at least a context or a sense in which ecological entropy H' can be thought of as “a *generalized measure of resistance to extinction*,” because this very same measure is then able to be repartitioned and reapplied *at every level* of organization or analysis as “*the*” stand-in proxy for the unlikelihood of extinction failure at that level. For each of the $(2^n - 1)$ conceivable family, geographic, or any other cross-sectional subgroupings of different potential cultivars, there is thus a

rigorous sense in which we may be entitled to conceptualize the design of efficient biomass proportions as involving essentially the same basic trade-offs and the same basic principles.

Corollary 1 indicates that the problem possesses enough structure to allow a striking characterization of efficient biomass proportions in the form of (52). The possibility of a tidy closed-form relation emerging from such a seemingly complicated problem is yet another unexpectedly neat consequence of the entropy expression H' appearing in the objective function.

Equation (52) implies that an optimal balance preserves *some* strictly positive level of biomass for *every* crop variety. Even if τ_j^* is *negative*, meaning preservation is costly for society, the model still wants *some* of crop variety j to be grown. The reason for such a strong conclusion has its origins based in the underlying model of endogenously determined disease. It is always worthwhile devoting *some* resources to preserving at least a small amount of each variety for a combination of the following two reasons. First of all, you never know when you may wish you had maintained the crop because you might want it some day if anything bad happens to some other crops. Taken alone, this is not nearly a sufficient reason for preservation, since it applies to almost everything on earth, and conservation is not free. But it becomes a powerful generic argument when combined with the second reason, as follows.

In this model, the “disease load” of endogenously determined host-specific pathogens on small host populations is so low that the correspondingly high host survivability *per unit of host biomass* represents, on the margin, a very attractive insurance investment. The model is abstracting away from all other sources of mortality, and hence, perhaps, emphasizes this limiting result excessively dramatically. Evaluated at near-zero biomass, in this world the marginal value of any species or variety is essentially infinite, because so few host-specific pathogens will be endogenously generated that a *relatively* significant enhancement of ecosystem survivability is possible on the margin. Therefore, the model wants us to devote *some* positive effort to preserving some small population of crop variety j , no matter how negative is τ_j^* , because it is still worthwhile *on the margin*. Implicit in such a viewpoint is an underlying dynamic perspective indicating that even one surviving acre of crop j may be valuable for rebuilding a j -based agriculture—if all else has failed. For even those least-promising crops having highly negative values of τ_j^* , the model is

willing to sacrifice (part of) an acre to j , but does not really care whether we actually plant the acre or just store some viable seed, provided merely that the equivalent opportunity-cost effort is expended to preserve the small acre's-worth of population.

The basic point that it is appropriate to maintain specimens of every variety is sufficiently central that some, possibly subdued, form of it should emerge, *mutatis mutandis*, even under a more general formulation. I think of the probability formulation of the paper as representing a no-frills core version of extinction by endogenously generated parasites, which, in the end, comes down to pathogen-induced deaths being randomly distributed among remaining host members at some quasi-fixed limiting rate, here ϵ . Any reasonable model of disease and epidemics, it seems to me, will be reduced, in the limit as the host population is approaching extinction, to the same basic mathematical structure involved here—although other such models may emphasize dynamic aspects of contagion omitted in this paper, or might be more disaggregated and more detailed, or may be phrased in the plant-genetics language of plague-resistant host phenotypes, and so forth.¹⁰

Thus, I believe the single most enduring message of this paper is that a relatively cheap way of buying catastrophe insurance is to cultivate or hold small positive amounts of as many different kinds of potential domesticates as it may be possible to preserve. The actual mathematical argument justifying such a conclusion, however, involves taking, and interpreting, some quite tricky limits, and goes to show the value here, indeed the necessity, of a rigorously formulated model.

It is interesting to note what happens to equation (52) for extreme values of λ . As λ is increased, the optimal distribution of biomass proportions becomes more evenly spread out among all of the varieties, approaching a uniform biomass fraction of $1/n$ per crop in the limit as $\lambda \rightarrow \infty$. As λ is made smaller, the solution becomes increasingly concentrated on successively narrower subsets of only the more profitable "higher- r " crop varieties. In the extreme limit as $\lambda \rightarrow 0$, there is complete specialization to the crops having the highest profitability. In the purely private $\lambda = 0$ competitive equilibrium corresponding to (13), the only crops

10. See, e.g., Burdon [1987]. Sometimes the language and jargon employed by plant pathologists is so specific as to give the misleading impression that the underlying principles are other than basic ecological concepts applicable to host-parasite-pathogen relations generally.

grown commercially will each have *zero* net marginal economic profitability—but there will also be a very large residual subset of potentially usable crop varieties with *negative* net marginal economic profitability, which stand for wild or obsolete strains that are ignored, or even eradicated to make way for more economically profitable land usage.

Two opposing forces are at play here. An “economic force,” represented by *low* values of λ , pulls the system toward specialization on “higher- τ ” varieties, which means severely unequal biomass proportions compared with any natural distribution. An “ecological force,” symbolized by *high* values of λ , pulls the system toward more equalitarian biomass proportions, stylistically closer to the way nature arranges things. For less extreme values of λ than zero or infinity, efficient combinations reflect a balance between these two opposing forces.

I will refrain here from inflicting on the reader a detailed exposition of everything that is left *out* of the model, since such an endeavor might well constitute a paper of its own. The core issue, as usual, is whether or not there is some essential misrepresentation of reality—something critical left out that should be included or something critical put in that does not belong—which might undermine the framework or the conclusions. Here I will give but two examples of “tweaking” the model with alternative formulations.

As a first example here, consider the idea that *positive* externalities may be generated by a farmer planting more of the same crop. One could tell stories about encouraging the development of complementary technologies because biotech companies are more likely to develop new herbicides, pesticides, seeds, fertilizers, and so forth for crops that are more widely planted.

I do not think that such type of example will alter the fundamental conclusions of the paper. First of all, absent some underlying economy of scale, the example appears to involve a pecuniary externality. The issue that the paper concerns, with or without genuine biotech externalities, is whether it is more economical for society as a whole not to leave quite so many flammable materials lying around, instead of having to extinguish the consequent fires when they inevitably arise. If we are talking about genuine nonpecuniary positive externalities, then these are unlikely to disrupt the main conclusion that it is appropriate to maintain some specimens of every variety, because the marginal diversity benefit goes to infinity for small crop sizes, while any

nonpecuniary positive externalities to growing more of an existing crop are likely to be quite bounded.

As a second example here, consider what might happen if, as is undoubtedly true in reality, the probabilities going into the calculations of the paper are not actually independent. A pathogen wiping out one strain or stand of a crop is more likely to wipe out a closely related strain or a closely located stand. In this case, the micro-level individual-crop equations corresponding to (52) and (53) will be modified by the presence of various combinations of relevant correlation coefficients, and some complicated generalization of the entropy formula will be required. But, as far as I can see, the big-picture macro-level view will be unaffected.

CONCLUSION

In the final analysis, what is one to take away from this model, which is expressed at such a very high level of abstraction? My hope is that such a way of combining economic with ecological modes of reasoning may be seen as giving some useful basic insights generally, and perhaps, for this particular issue of optimal cropping proportions, more useful basic insights than can be gotten from either worldview alone.

HARVARD UNIVERSITY

REFERENCES

- Begon, M., J. L. Harper, and C. R. Townsend, *Ecology*, 2nd edition (London: Blackwell Scientific Publications, 1990).
- Burdon, J. J., *Diseases and Plant Population Biology* (Cambridge, UK: Cambridge University Press, 1987).
- Connor, E. F., and E. D. McCoy, "The Statistics and Biology of the Species-Area Relationship," *American Naturalist*, XIII (1979), 791–833.
- Cowell, Frank A., *Measuring Inequality*, 2nd edition (London: Prentice Hall, 1995).
- Huston, M. A., *Biological Diversity* (Cambridge, UK: Cambridge University Press, 1994).
- MacArthur, Robert H., and Edward Wilson, *The Theory of Island Biogeography* (Princeton, NJ: Princeton University Press, 1967).
- Meffe, G. K., and C. R. Carroll, *Principles of Conservation Biology* (Sunderland, MA: Sinauer Associates Inc., 1994).
- Pianka, E. R., *Evolutionary Ecology*, 5th edition (New York: Harper and Row, 1995).
- Pielou, E. C., *Mathematical Ecology*, 2nd edition (New York: John Wiley and Sons, 1977).
- Raisbeck, G., *Information Theory: An Introduction for Scientists and Engineers* (Cambridge, MA: The MIT Press, 1964).
- Schwarz, Michael, "Decision Making under Extreme Uncertainty," mimeo, Stanford University Graduate School of Business, November 30, 1998.
- Sears, Francis W., *Thermodynamics, the Kinetic Theory of Gases, and Statistical Mechanics*, 2nd edition (Reading, MA: Addison-Wesley Publishing Company, Inc., 1953).
- Theil, Henri, *Economics and Information Theory* (Amsterdam: North-Holland Publishing Co., 1967).

This article has been cited by:

1. Jouni Korhonen, Thomas P. Seager. 2008. Beyond eco-efficiency: a resilience perspective. *Business Strategy and the Environment* 17:7, 411-419. [[CrossRef](#)]
2. Françoise Benhamou, Stéphanie Peltier. 2007. How should cultural diversity be measured? An application using the French publishing industry. *Journal of Cultural Economics* 31:2, 85-107. [[CrossRef](#)]
3. Jouni Korhonen. 2006. On the paradox of corporate social responsibility: how can we use social science and natural science for a new vision?. *Business Ethics: A European Review* 15:2, 200-214. [[CrossRef](#)]
4. Francois Moreau, Stephanie Peltier. 2004. Cultural Diversity in the Movie Industry: A Cross-National Study. *Journal of Media Economics* 17:2, 123-143. [[CrossRef](#)]
5. Timo Goeschl , Timothy Swanson . 2003. Pests, Plagues, and PatentsPests, Plagues, and Patents. *Journal of the European Economic Association* 1:2-3, 561-575. [[Abstract](#)] [[PDF](#)] [[PDF Plus](#)]
6. Jouni Korhonen. 2002. The dominant economics paradigm and corporate social responsibility. *Corporate Social Responsibility and Environmental Management* 9:1, 66-79. [[CrossRef](#)]