CHAPTER 12

Market Failures and Biotic Resources

With the exception of solar energy and freshwater in the solar-powered hydrologic cycle, abiotic resources are nonrenewable on human time scales. Their value to humans consists solely in their extraction and use. Biotic resources, on the other hand, are renewable, if degradable, and are valuable as much for the services they provide as for any goods that can be derived from them. In this chapter, we examine whether or not specific natural resources meet the criteria for market allocation, turning our attention to biotic stock-flow resources (ecosystem structure) and ecological fund-service resources (ecosystem function), paying special attention to waste absorption capacity.

RENEWABLE RESOURCE STOCKS AND FLOWS

Renewable resource stocks and flows are rival and potentially excludable, depending on whether or not institutions exist that can regulate access to them. If depleted at a rate no faster than they regenerate, they are nonrival between generations. Unfortunately, unless we explicitly take future generations into account, economic incentives quite likely will lead us to deplete many of these resources faster than they can regenerate, and may eventually threaten them with extinction. What's more, as we have pointed out repeatedly, the use of renewable resource stocks and flows unavoidably deplete ecosystem funds and services as an "externality" of their production. This dramatically complicates economic analysis of these resources.

In our earlier discussion of renewable resource stocks and flows, we looked at their physical properties with little if any discussion of the economics involved. Recall the sustainable yield curve from Chapter 6, here reproduced as Figure 12.1.[1] At first glance, it would appear that the goal of economists would simply be to make the resource as productive as possible. If this were the case, we should strive to maintain a population that produces the maximum sustainable yield, or MSY. However, this ignores two major issues. First, there are costs to harvesting, which we will call $P_E E$ (price of effort times effort, where effort includes all the resources required to harvest a stock), and these costs are likely to increase per unit harvested as the population in question grows smaller. Obviously, the smaller the population of fish that remains, the harder they are to catch. Even for forests, the most accessible timber will be harvested first, and as forest stocks decrease, it will cost more to bring the less accessible stocks to market.

Second, and even more important, if we were to consider all resources substitutable, as many economists do, and money is the perfect substitute for any resource, then the economic goal would not be to maximize the sustainable harvest of any specific resource, but rather the monetary sum of annual profits yielded by the resource. But even that is incomplete, as we will discuss shortly, for the market goal is in fact to maximize present value—the monetary sum of discounted future profits.

To simplify analysis, we can assume a linear relationship between effort, stock, and harvest, known as the **catch-per-unit-effort hypothesis**. For any given effort, more stock leads to a larger harvest in a linear fashion, and for any given stock, more effort yields a larger harvest. However, unless harvests are less than the annual increase in stock (i.e., below the sustainable yield curve), an increase in effort in any given year, all else being equal, implies a smaller stock and thus a smaller harvest from the same level of increased effort in the following years.[2]

Using simple algebra, we could assume that $Y = qXE$, where Y is harvest, X is stock, E is effort, and q is a constant we can think of as the "catchability coefficient."[3] As those of you at all familiar with math will recognize, this is the equation for a line starting at the origin with slope equal to effort, E. In Figure 12.1, we have drawn in lines $Y = qXE$, $Y = qXE'$ and $Y = qXE''$, where $E'' > E' > E$. Say the stock is Q″ in year 0, and effort is E. Our harvest then is Q. Of this harvest, Q″Q′ corresponds to the

---

[1] We changed the scale of the figure to make it clearer. Specifically, the Y-axis has been stretched, so that the 45-degree line appears to be steeper than 45 degrees.

[2] When population stocks are large, a large harvest can lead to a larger sustainable yield in the following year, but this sustainable yield will still be smaller than the harvest required to reduce the population to that stock that provides that sustainable yield.

[3] C. Clark, *Mathematical Bioeconomics: The Optimal Management of Renewable Resources*, New York: Wiley, 1990.

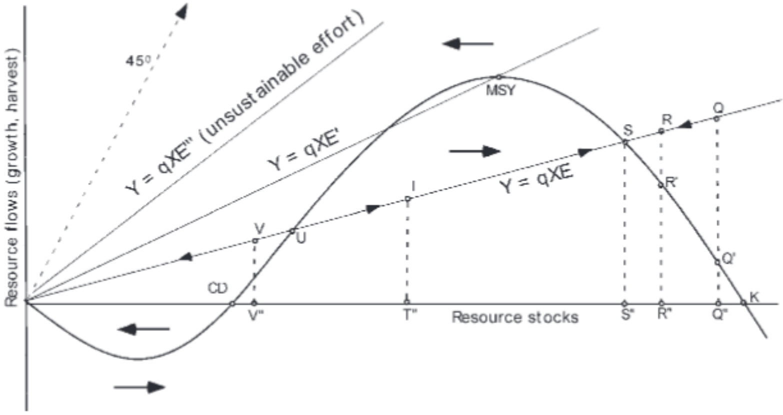

Figure 12.1 • The sustainable yield curve and catch-per-unit-effort curves. Recall from Chapter 6 that the curved line represents the rate of growth of stock for any level of stock, which is the same as sustainable yield. The straight lines Y = qXE represent the harvest for a given effort at any level of stock. A steeper slope indicates greater effort.

annual growth for that year, and Q′Q must therefore reduce the stock, to R″. (Note that this figure is not drawn with the same scale on both axes; i.e., resource flows are in smaller units than resource stocks.) At R″ and effort E, the harvest will be at point R, reducing the stock by R′R. As long as effort remains constant, this process continues until we reach stock S″ and point S on the sustainable yield curve.

Now imagine an external shock, such as El Niño, pushes populations down to T″ in a given year. As long as effort stays the same, the annual harvest will now be less than the growth increment or net recruitment, and the stock will gradually recover, until we again reach S″. S is therefore a *stable equilibrium point*. However, if another El Niño year occurs before the fish population has recovered, it could push the population down to V″. At V″, the same effort will lead to a harvest greater than the growth increment, and the population will not recover. Recall from Chapter 6 that any harvest below the sustainable yield curve leads to a higher stock in the following year, and any harvest above the curve leads to a lower stock. The arrows on the Y = qXE curve illustrate this dynamic. Thus, point U, where the catch-per-unit-effort curve intersects the sustainable yield curve, is an unstable equilibrium.

Maximizing Annual Profits

Suppose the goal is to maximize sustainable annual profits (π) from exploiting the fishery. This requires that we figure out where on the sustainable yield curve profits are maximized. Unfortunately our graphical analysis using the effort curve does not directly show profit. To analyze the question from the perspective of annual profit is worthwhile and will require a somewhat different graph (Figure 12.2). The axes and the yield curve remain the same, except that we multiply the vertical (flow) axis by an assumed constant, P_F, the price of fish. This converts the yield curve into a total revenue (TR) curve without changing its shape, since we are multiplying by a constant. Since profit, π, is equal to TR–TC (total cost), we need to add a TC curve. If we define effort as all the equipment, labor, and other resources that go into fishing, then TC is equal to the amount of effort times the price of effort, and therefore can be derived from a series of catch-per-unit-effort curves.[4] We may think of TC as a curve starting at maximum population and rising to the left as more fish are caught. TC will increase as more fish are caught both as a result of stock depletion and as a result of harvesting a larger sustainable yield (at least up to MSY). But even beyond MSY the TC will probably still rise because stocks have become so sparse that the fish have become hard to find.[5]

Focusing on the flow dimension only, π = TR–TC, and maximum π occurs at π^* where MR = MC, where the slope of the tangent to the TR curve (MR) equals the slope of the tangent to the TC curve (MC). Remember that we have left the stock dimension out of our concept of total revenue—we are analyzing profit as a sustainable flow, not as the result of unsustainable stock reduction. Some stock reduction is necessary to arrive at the profit-maximizing stock, but ignore that for now; just assume for now that the stock-reducing, one-time catch of fish are thrown away. We'll deal with them later.

---

[4] Technical note on the derivation of Figure 12.2 from Figure 12.1: The catch-per-unit-effort curve ($Y = qXE$ in Figure 12.1) tells us the amount we can harvest from a given level of effort (E) at any level of stock. The TC curve tells us how much it costs to achieve any sustainable harvest. Theoretically there are two levels of stock (one level of stock when there is no critical depensation) where the harvest from a given constant effort is sustainable, and they are given by the intersection of the $Y = qXE$ curve with the sustainable yield curve. But in a world full of stochastic shocks and constant change, the unstable equilibrium will not be sustainable and can be ignored. The Y coordinate of the TC curve is therefore E^*E_P (where E_P = the constant price of effort), and the X coordinate is given by the stock at which the yield from a given $Y = qXE$ curve is sustainable, i.e., the stock at which $Y = qXE$ intercepts the sustainable yield curve. By letting E increase from zero to maximum sustainable effort, one can map out the TC curve. Total cost will be zero at K but rises until it reaches the maximum sustainable level of effort. Greater effort increases the sustainable yield until MSY is reached, at which point additional increases in effort lead to smaller and smaller sustainable yields. Eventually, increases in effort become unsustainable. At this point, the TC curve is truncated, and the dashed line in Figure 12.2 indicates unsustainable levels of effort. With the sustainable yield and constant effort curves of the type used in Figure 12.1, the TC curve in Figure 12.2 would not be a straight line; but a straight line facilitates the analysis and is no greater an abstraction from reality than any of the other curves depicted.

[5] Note that Figure 12.2 is not the same shape as Figure 12.1. Figure 12.2 depicts a more resilient population that can sustain higher levels of effort without causing extinction, to better illustrate the points we are making here.

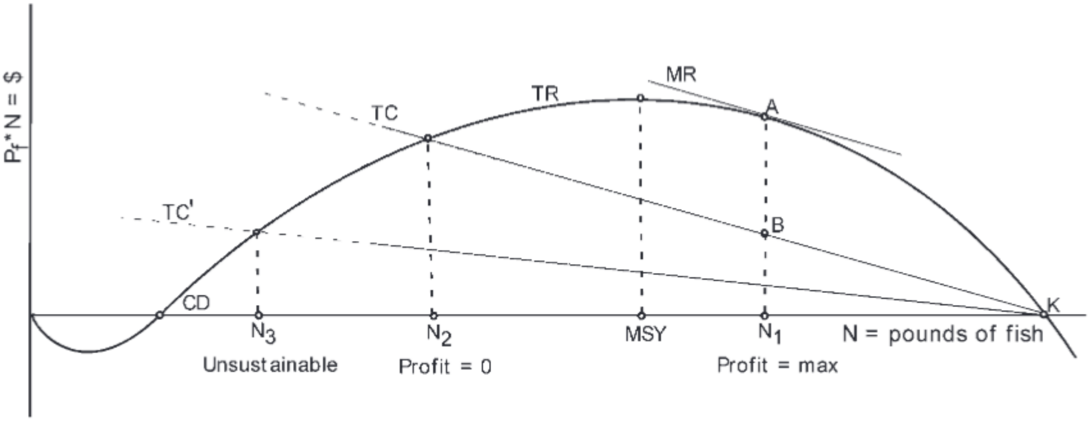

Figure 12.2 • Maximizing annual profit from renewable resources.

In Figure 12.2, we can see that the maximization of annual profit (AB) occurs at N_1, which is larger than the stock corresponding to MSY. In other words, in this analysis the profit-maximizing capitalist will not even reach MSY, much less drive the population to extinction! If total costs were zero (or constant), marginal costs would be zero, and profit maximization would occur where marginal revenue was also zero, namely at MSY where the tangent to the TR curve is horizontal (slope equals zero). So even zero harvest cost would not lead the capitalist to exploit beyond MSY!

THINK ABOUT IT!

Can you find the tacit assumption responsible for the happy result that profit-maximizing exploitation does not require much stock depletion?

We have assumed a single capitalist exploiting the fishery, a single owner or decision maker. Instead of private property with excludability, suppose the fishery were open access, as most are. Under an open access regime (in which you'll recall the resource is nonexcludable), new fishermen will enter as long as there are profits to be made.[6] New entrants will push the stock down to N_2 at which π = zero, or (TR = TC). At N_2 many more resources are going into fishing, but the sustainable catch is less than at N_1, and no one is making a profit.[7]

Curve TC′depicts lower harvest costs, which might come about from a

---

[6] Profits, in economic parlance, are returns above and beyond the cost of production, where the cost of production obviously includes wages. For a small fishing crew sharing returns, profits would mean higher wages than they could find elsewhere.

[7] Note that if total costs are high enough, it is possible that the open access equilibrium might be at a greater yield than the profit maximizing equilibrium, though still with zero profit.

technological advance such as sonar devices for locating fish schools. At these lower costs, in an open access fishery it would be profitable for new fishermen to keep entering the fishery even after harvests become unsustainable. This is the case where the tragedy of open access resources may well lead to extinction, and may be a realistic depiction of what was happening with North Atlantic cod and many whale populations before regulation began.

Box 12-1 Annual Profit Maximization in Wild vs. Bred Populations

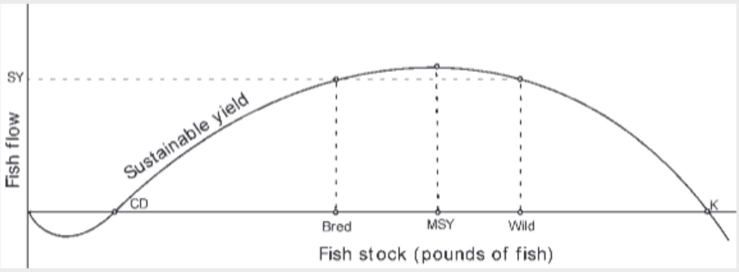

Figure 12.3 • Sustainable harvests from wild vs. bred populations.

Before leaving the annual profit-maximizing approach, we'll use it to distinguish between the exploitation of wild and bred populations. The fisheries example we have used is, of course, the case of a wild population. A catfish pond would be an example of a bred population. In both cases, the biological population growth function is similar. But the cost functions are very different. For a wild population, costs are mainly costs of capture. For a bred population, costs of capture are minimal but the costs of feeding and confinement are high.

Figure 12.3 shows two equal sustainable yields on a given population growth curve, one corresponding to a wild mode of exploitation and the other a breeding mode. In the wild mode, the TC curve (not shown) would start from the right and rise to the left (increasing with capture, as in our fisheries example). In the breeding mode, the TC curve would start from the left and rise to the right (increasing with feeding and confinement costs of the larger population). The graph depicts a case in which the same SY could be had from exploiting the population in either a wild or a bred mode. Note that just as the annual-profit-maximizing owner of a wild fishery will never push the population below MSY, the profit-maximizing owner of a bred fishery will never allow the population to grow up to the point of MSY. Can you explain why?

The advantage of the wild mode is that the base population is larger, containing more biodiversity, providing more ecological fund services, and feeding is free. The advantage of the bred mode is that the smaller base population takes up less ecological space, leaving more room for other life, wild or human. Of course the food for the captive population requires some ecological space as well as collection. Controlled breeding and genetic engineering would seem to require the breeding mode. Burgeoning human populations and the technical thrust toward genetic engineering push the shift toward bred genetically engineered populations (e.g., the famous 200-pound salmon). Also, the quest for higher biological growth rates to keep up with the interest rate favors the bred stock, since a smaller base population yielding the same annual increment obviously implies a higher rate of growth. But smaller populations growing more rapidly, and ever more dependent on human management for feeding, reproduction, and disease control, surely will increase the instability, brittleness, and vulnerability to uncertainty of the whole biotic system. Do you think the Federal Reserve Board considers these factors when setting interest rates? Should this risk be factored into interest rates, and by extension, discount rates?

Profit-Maximizing Harvest When Profits Can be Invested: Net Present Value

Returning to the higher-cost scenario (TC), suppose we make sure that there is a single owner, not open access. Are we then sure that we will end up at N_1? Unfortunately not, because of the troublesome issue we swept under the rug initially: What happens to those fish that are part of the stock reduction rather than the annual recruitment? They are not thrown away, and their number is large relative to the annual growth. Those fish are sold. Stock reduction fish have the advantage of being available now—you don't have to wait for them to be hatched and grow. But the more you reduce the stock of fish today, the fewer fish you will have tomorrow, and the more difficult it will be to catch those fish. The population of fish is like the proverbial goose that lays golden eggs in perpetuity. Surely no rational capitalist would kill such a productive goose.

Or would she?

If the capitalist wanted to maximize the sum of golden eggs from now till the end of time, then obviously she would not kill the goose. But the goose also has a liquidation value as a cooked goose. Suppose the capitalist could kill the goose, cook it and sell it for a sum of money, which when put in the bank at the going interest rate would yield an annual sum greater than the value of the golden eggs? Then it's goodbye goose, hello bank! The population growth rate of the goose (its egg-producing fecundity) is in direct competition with the interest rate, the "fecundity" of money. Neoclassical economists argue that money itself may have no reproductive organs, but it is a surrogate for many other things that can reproduce, and on average those other things can reproduce faster than the goose. So the goose-killing, reinvesting capitalist has converted a slow-growing asset into a fast-growing one, and we are all therefore better off. According to economists, cooking the goose in this case maximizes **net present value (NPV)**—the value to us today of all cost and benefit streams from now into the future. Economists calculate NPV by using a discount factor to give less weight to costs and benefits the farther in the future they occur (see Chapter 10).

Let's take the story a bit further in a thought experiment. Suppose an economy consists only of renewable resources. The interest rate is equal to some weighted average of the growth rates of all renewable resource populations. Everything that grows more slowly than the average (the interest rate) is a candidate for extinction (unless at some stock its growth rate rises above the interest rate). But something is always below average. When the below average is eliminated, what happens to the average in the next period? It goes up, of course. The tendency, it seems, would be to end up with only the fastest-growing species. Biodiversity would entirely disappear. In a world in which everything is fungible,[8] that would not matter. We could all eat algae, if that were the fastest-growing species.

But we have forgotten prices. Surely prices would rise as particular slow-growing species became scarce and the rising price would compensate for low biological growth rate, so that the value of the species would grow at a rate equal to the rate of interest before it became extinct. Yes, but remember that when the price goes up, the price of the existing stock rises as well as the price of the flow of recruits. As the price increases, the incentive to liquidate the now more valuable remaining stock rises, along with the incentive to reduce current offtake to allow an increase in the more valuable new recruits.

The bluefin tuna is an excellent example of this argument. In 2001, a single 444-pound bluefin tuna sold in Japan for nearly $175,000, or about $395/lb. Although this was an anomaly, restaurants in Japan regularly pay up to $110 per pound for bluefin tuna.[9] Admittedly this occurs

---

[8] Something is fungible if one unit of it substitutes indifferently for another unit. For example, two buckets of water from the same well are fungible (you can't tell any difference between them). But two buckets of water from two different wells may not be fungible because of qualitative differences such as hardness, taste, etc. Money is fungible; we cannot tell if the money the government spends on foreign aid came from my tax dollars or from yours. Things convertible into money—goods, services, even biological species—acquire a kind of artificial or abstract fungibility, even though physically they are not at all fungible. This makes it easy to commit the fallacy of misplaced concreteness (see Chapter 2).

[9] G. Schaeffer, Tuna Sells for Record $175,000, Associated Press International, January 5, 2001.

under a regime of imperfect property rights,[10] but how confident are you that private ownership of the bluefin would solve the problem? The higher price means higher liquidation value, as well as higher future revenue, from new tuna.

How do we decide whether or not to harvest a (marginal) ton of fish that if left in the water reproduces (giving us a flow of golden eggs) and if harvested (giving us a cooked goose) yields interest on the profit?

The neoclassical approach is to ask: What are all the opportunity costs of harvesting the fish today? Obviously, if we harvest the fish today, that same ton of fish is not there to harvest tomorrow. Unlike oil or iron, however, we lose not only the opportunity to harvest that ton of fish, but also the offspring those fish would have had, as well as the increase in biomass they would have experienced, if left in the water. In addition, as the economy grows, people demand more fish, and growing human populations further increase the demand for fish. Greater demand means a higher price, and we lose that additional profit if we harvest today. Moreover, if we leave the fish in the water, they will reproduce, and the greater population means that it will be cheaper and hence more profitable to catch a ton of fish next period than this period.

In contrast, the benefit of catching the fish this period rather than next is that the profits from their harvest can be invested; that is, the opportunity cost of not catching the fish is the money forgone from not being able to invest the profits from that ton of fish between this period and the next. As we harvest more and more fish, the growth rate (the annual increase in biomass/total biomass) increases, so the marginal costs of harvesting now are increasing.

The economist will therefore favor harvesting as long as the marginal benefits of catching the next ton of fish are greater than the marginal costs, and will stop when they are equal. The tricky thing about the decision is that when we consume more stock today in exchange for less stock and less yield tomorrow, it is not just one tomorrow, but tomorrow and tomorrow and tomorrow in perpetuity. We must compare a one-time benefit with a perpetual loss. As we mentioned above, economists address this problem rather unsatisfactorily by the financial convention of discounting and present value maximization. What does all of this mean in practical terms?

Compared to our static analysis when we ignored the stock reduction

---

[10] Fishing of bluefin tuna in the Atlantic at least is regulated, with quotas for the Eastern Atlantic held by European countries, and quotas in the Western Atlantic held by the U.S., Canada, and Japan. But evidence indicates that these quotas are too high, and that the two populations are not even distinct. Thus, regulations designed to assign property rights to a formerly open access resource may fail in their objective of preserving the species. See T. Bestor, How Sushi Went Global, *Foreign Policy* 121 (November/December 2000).

required to reach the annual profit-maximizing equilibrium, the opportunity to invest profits from stock reduction will lead to a lower stock of fish (or any other renewable resource). In terms of Figure 12.2, the profit-maximizing harvest will be to the left of N_1, and if the interest rate is high enough relative to the growth rate of the species, it may even be to the left of MSY.

Take the case of bluefin tuna, where the cost of capture of one fish may be a negligible portion of its sales price. Imagine current harvests were sustainable and in the vicinity of MSY, and there was a single resource owner intent on maximizing profits. The reduction in sustainable yield from MSY to a somewhat lower stock may be small, while the stock liquidation required to get there is still quite large. If the interest payments on the profits from the sale of that liquidated stock are greater than the value of the lost annual yield, then profit maximization favors sustainable harvest at a stock lower than MSY, and closer to the point of critical depensation.

What are the implications of this scenario in an extreme but not at all unrealistic case where harvest costs (P_EE) are negligible compared to harvest revenue (P_YY), even for very low-resource stocks? Timber is a good example. Say that a forest of redwoods not yet biologically mature[11] increases in size and thus value by 3.5% per year.[12] In contrast, the average real growth rate of money on the U.S. stock market over the last 70 years was about 7%.[13] Clearly, it makes economic sense to harvest the resource now and invest the profits in the bank. In fact, for any species that is relatively inexpensive to harvest and grows more slowly than alternative investments, it makes economic sense to harvest the species to extinction. In general, averaged over the time it takes to reach harvest size, many valuable species grow quite slowly relative to alternative investments, and technology tends to reduce unit harvest costs over time. And for such resources it is, once again, goodbye golden eggs, hello bank.

In summary, the advantage of the catch-per-unit-effort curve analysis is that it builds in from the beginning the stock reduction effect and in that way is more realistic. The advantage of the TR–TC diagram is that it shows

---

[11] Biological maturity occurs when growth rates for the forest taper off toward zero; i.e., new growth is just matched by rates of decay.

[12] This is actually an unrealistically high rate of growth. Data exists on a 1-acre plot of redwoods that has been monitored for over 70 years. Though the rate of growth on this plot is so high that it is widely known as the "wonder plot," the most rapid 10-year mean annual increment in total stand volume was 3.5%. The mean annual increment in growth from 1923 to 1995 is well under 1%. These figures were calculated by the authors from data provided in G. Allen, J. Lindquist, J. Melo, and J. Stuart, Seventy-Two Years Growth on a Redwood Sample Plot: The Wonder Plot Revisited (no date); Online: http://www.cnr.berkeley.edu/~jleblanc/WWW/Redwood/rdwd-Seventy-.html.

[13] S. Johnson, Are Seven Percent Returns Realistic? Online: http://www.sscommonsense.org/page04.html. Common Sense on Social Security.

that annual profit maximization can be sustainable and efficient in the absence of open access, and with some limits on the biologically blind financial logic of discounting and present value maximization.

Renewable Resource Funds and Services

The analysis of optimal harvests of renewable resources so far has only treated them as stocks and flows of raw materials. But, as we discussed in Chapter 6, renewable resources are also funds that provide ecosystem services, and we cannot ignore one when deciding how best to allocate the other. While natural resource stocks and flows have some characteristics of market goods, the services generated by funds typically do not. Such services are generally nonexcludable and for many, no feasible institutions or technologies could make them excludable. Thus, free markets will not produce them. They are also nonrival and noncongestible, and selling them in the market would not equate marginal costs with marginal benefits.

Treating the destruction of ecosystem services as a negative externality of aggregate economic production offers some insights into optimal harvest levels. Returning to the analysis of renewable resource stock-flows, we would need to add all external costs to total private harvest costs. If marginal external costs increased linearly as the fund-stock is reduced in size, then total external costs would increase exponentially. In reality, marginal external costs are likely to increase at a greater than linear rate as we near an ecological threshold, such as the critical depensation point, which will inevitably be reached when harvest effort is too great to be sustained. Figure 12.4 is similar to Figure 12.2 but relabels the TC curve as the total private cost (TPC) curve and includes a total social cost (TSC) curve that adds external costs to the TPC. The TSC curve approaches the vertical

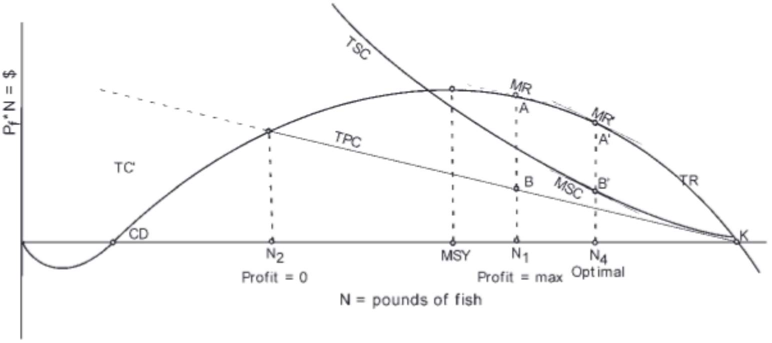

Figure 12.4 • Optimal harvest of renewable resources when accounting for ecosystem services.

(i.e., unacceptably high marginal social costs) as we approach the point of critical depensation. The optimal harvest is where marginal social costs equal marginal revenue, labeled N_4 on the graph. Whenever the renewable resource contributes to the provision of ecosystem services, the optimal harvest from an ecological economic perspective will always be at a higher stock with lower private costs than in the annual profit-maximizing equilibrium.

Of course, optimality would require the micro-level internalization into prices of all ecosystem services. Yet human impacts on these services are characterized at best by uncertainty (we know the possible outcomes of damage to ecosystem funds on ecosystem services, but don't know the probabilities) and more often than not by ignorance (we don't even know the range of possible outcomes). In fact, we almost certainly do not know the full extent of the ecosystem services from which we benefit. In addition, the value of all externalities would need to be worked out by economists, ecologists, and others, and incorporated into the prices of the goods that generate the externalities. And, of course, the marginal value of an ecosystem service changes along with the quantity of the ecosystem service supplied, so the value of externalities would be constantly changing. As we have pointed out, all economic production incurs externalities. The notion of calculating the constantly changing values of all externalities for all goods would be a Promethean task. Once achieved, it would still require some institution to incorporate the fees into market prices. And we must remember that the magic of the market is precisely its unplanned, decentralized nature, and its ability to utilize "knowledge not given to anyone in its totality."

Effectively internalizing externalities, in contrast, requires precisely the opposite—centralized planning by individuals provided with knowledge in its totality. While an optimal allocation of everything is not a feasible goal, in Part VI on policy (Chapters 20–23), we will explore approaches to achieving a satisfactory allocation.

Note also that we have again left out the potential for investing the profits from reducing the stock. This is intentional. No potential return could substitute for either life-sustaining ecosystem services or the raw materials essential for all economic activities. In addition, many investments are profitable precisely because they do not account for the opportunity costs of resource depletion (MUC) or the social costs of the ecosystem services inevitably degraded or destroyed through resource extraction.

Nonmarket ecosystem fund-services simply cannot be converted to money and invested, like we can do with a cooked goose. Also, natural resources are growing physically scarcer. Technology seems to be developing new uses for most natural resources faster than it develops substitutes, which increases future demand. Increasing demand and decreasing sup-

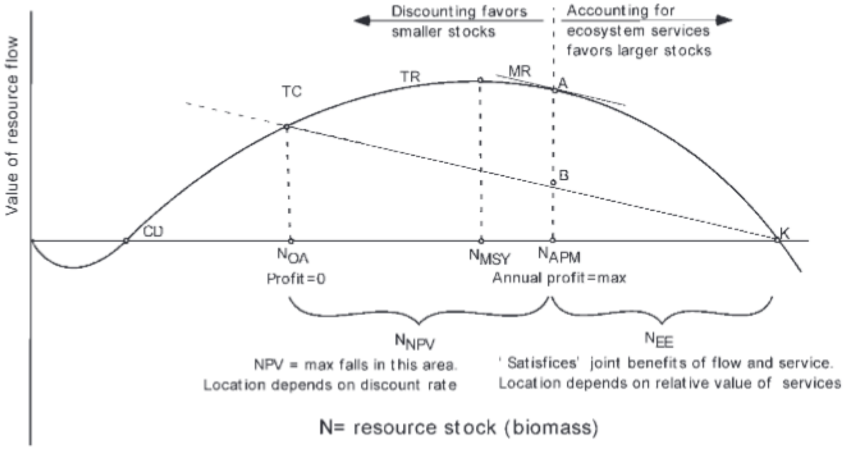

Figure 12.5 • Optimal harvest levels for renewable resources with respect to different objectives and management regimes. From left to right, N_{OA} is the open access equilibrium, at which profit is zero. If total costs decrease over time, the N_{OA} may become unsustainable. N_{NPV} is the stock at which net present value is maximized. At very high discount rates, this will be the same as the open access equilibrium, and at a zero discount rate will be equal to N_{APM}. N_{APM} is the annual profit-maximizing stock. N_{EE} is the objective of ecological economists, and strives for "satisficing" (seeking a sufficient, rather than the maximum, amount) the joint benefits of both flow and service.

ply imply greater value for natural resources in the future, not less. Instead of intertemporal profit maximization, we concur with Geoffrey Heal and other environmental economists that we should seek to maximize well-being from renewable natural resources for the current generation without diminishing the capacity of future generations to benefit from those resources. Heal and others have called this principle the Green Golden Rule.[14] Applied to the stock-flow alone, the Green Golden Rule corresponds to our analysis of maximizing sustainable annual profit, as opposed to present value maximization.

In summary, the more we have of an ecological fund, all else being equal, the more services we can expect it to provide. If we are concerned only with the service provided by a fund, the optimal amount of the fund is the carrying capacity, as measured on the X-axis in Figures 12.1 and 12.4. In contrast, optimal harvest of stocks is solely a function of flow (the Y-axis in the same figures). Unless we recognize the values of both the fund-service and the stock-flow, resource extraction rates will not be optimal. Figure 12.5 summarizes the discussion of optimal stocks and harvests from renewable resources.

---

[14] G. Heal, *Valuing the Future: Economic Theory and Sustainability,* New York: Columbia University Press, 1998. Note that Heal proposes a number of objective functions in addition to the Green Golden Rule.

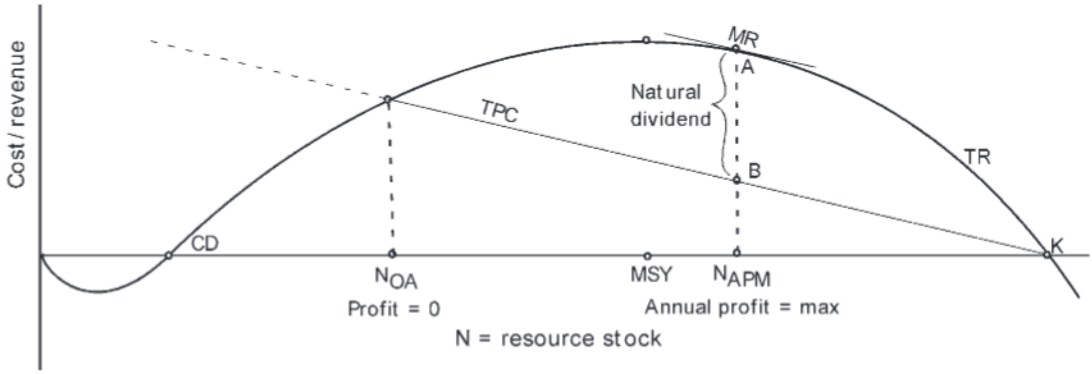

Figure 12.6 • The natural dividend from renewable resource harvests.

The Natural Dividend From Renewable Resources

Unearned income is the amount above and beyond what is required to bring a resource to market. In the case of nonrenewables, the unearned income was called scarcity rent. In the case of renewables, there can also be an unearned income deriving from nature's reproductive capacity, which we propose to call the **natural dividend**. Extraction and harvesting impose real costs, and those engaged in these activities are earning a legitimate income. This is included in the TC curve as "normal profit," which is the opportunity cost of the owner's labor and capital. Profit above normal profit (e.g., AB) is called "pure economic profit." Since it is beyond the opportunity cost of the owner, we think of it as an unearned growth dividend from the reproductive power of nature. We have already seen that in an open access equilibrium, total costs are just equal to total revenues, and profits are zero (at stock N_{OA} in Figure 12.6). However, if there is a single owner of the resource, say, a government agency, the owner can limit harvest to the profit maximizing level, stock N_1. The pure profits in this case (AB in Figure 12.6) are the natural dividend and arise not from special abilities of the particular owner, but from the reproductive powers of nature. The natural dividend can also be thought of as the value of the resource in the ground (or sea). While the natural dividend typically accrues to the owner of the resource, it is a purely unearned profit from production. To whom should the dividend belong? That is a political decision, more in the realm of fair distribution than efficient allocation.

The *natural dividend* is the unearned income from the harvest of renewable resources. As nature and not human industry produces renewable resources, all profits above "normal" profit (included in TC) are unearned, and the natural dividend is equivalent to TR–TC.

WASTE ABSORPTION CAPACITY

Waste absorption capacity is really just another ecosystem service. We treat it separately here because it is extremely important and because it has different characteristics from most other ecosystem services. Waste absorption capacity is the ability of the ecosystem to absorb and process pollution, and the economics of pollution is the predominant focus of neoclassical environmental economics. As we pointed out previously, waste absorption is a rival good. If I dispose of my sewage in a wetland, there is less capacity subsequently for that wetland to process someone else's wastes. As we also pointed out, many countries are trying to create institutions that make waste absorption capacity an excludable good. These can range from regulations that directly limit industrial emissions to mandatory catalytic converters in cars to tradable emissions permits for sulfur oxides. Tradable permits and quotas for pollution essentially make waste absorption capacity a private good. These mechanisms will be discussed at length in Chapter 21.

However, we must bear in mind that pollution is pure externality and a **public bad**, which is something that is nonrival, nonexcludable, and undesirable. Therefore, even when a market exists in air pollution, for example, this is not the same as saying that a market exists for clean air. Nonetheless, pollution permits are one of several mechanisms that can help achieve a socially optimal level of pollution (see Chapters 20–22).

There is, of course, no direct social benefit to pollution *per se*, but as we have repeatedly stated, all productive processes generate some pollution, and if we prohibit all pollution, we virtually prohibit production. This is why economists use the apparent oxymoron of "optimal pollution." By optimal, economists simply mean potentially Pareto efficient. A reasonable estimate of the benefits of pollution is therefore the marginal net private benefits of production (MNPB) associated with a unit of pollution. The problem is, of course, that our knowledge of external costs of pollution is characterized predominantly by ignorance and uncertainty, to a lesser extent by risk, and to a minimal extent by certainty. Since we do not know the full social costs of pollution, it is exceedingly difficult to balance costs with benefits. Policy makers are also not well informed concerning the MNPB of pollution to polluters.

We have to recognize that waste absorption capacity is a dynamic process, and we must define carefully what we mean by it. We define **waste absorption capacity** as the ability of an ecosystem to assimilate a given flow of waste. If the waste flow exceeds the waste absorption capacity, then waste will accumulate. As waste accumulates, the ecosystem becomes less able to assimilate it, leading to even more rapid accumulation. Thus, in a dynamic analysis, if the flow of waste into an ecosystem exceeds the ability of the ecosystem to assimilate that waste, then waste will build up indefinitely until virtually all services generated by the biotic processes in that ecosystem collapse. The system will thereafter be able to process only negligible amounts of waste.

What is the marginal external cost of the increase in flow rate that takes the system beyond the point of no return? It is the value of the lost services from that ecosystem for all time. If the ecosystem provides vital functions, the marginal external cost is basically infinite. However, if the waste flow is halted before collapse of the ecosystem, the system can slowly process the waste and restore itself. This is basically what is happening in Lake Erie on the U.S.–Canada border. If the ecosystem in question is simply a local ecosystem with similar ecosystems nearby, even after collapse, stopping the waste flow can lead to recovery. Wastes will be absorbed, dissipate, and settle out of the system, new organisms will colonize the system, and the restoration process will begin.

It is worth illustrating these points graphically. Figure 12.7 is an appropriately modified version of the analysis of externalities from Figure 10.2. Note that economic output and waste output are measured on the same axis in recognition of the laws of thermodynamics. In reality, the relationship is not as fixed as Figure 12.7 indicates; many technologies are available for producing different goods, some of which generate less pollution than others, though for any given technology, the relationship will be fixed.

We assume here that the economic output under consideration is not essential—either substitute products and processes are readily available, or the good itself is simply not that important to human quality of life. Chlorine-bleached paper from wood pulp is a good example. Unbleached paper from kenaf or hemp are excellent substitutes, and paper itself is important, but not essential to life. Among the many wastes emitted by paper mills are organochlorines, and paper mills are the largest emitter of organochlorines into the water supply in the U.S. and numerous other countries as well.

Organochlorines resist biodegradation, and hence the waste absorption capacity for these substances is quite small; it is indicated by the perpendicular at point Q_A, W_A in Figure 12.7. Organochlorines include some of the most toxic substances known, such as dioxins, which readily accumulate in the environment and in animals, including humans. Health problems associated with dioxins include cancer, immune system disorders, and developmental problems in children and fetuses. We assume without too much exaggeration that surpassing the waste absorption capacity for organochlorines for extended periods could make the affected areas essentially uninhabitable for humans, and on our crowded planet, this is an unacceptably high cost. This is indicated on the graph by the marginal external cost (MEC) curve, which approaches vertical as we near the waste absorption capacity of the environment (Q_AW_A). These nearly infinite marginal costs occur only if the accumulation of pulp mill waste continues unabated for some time.

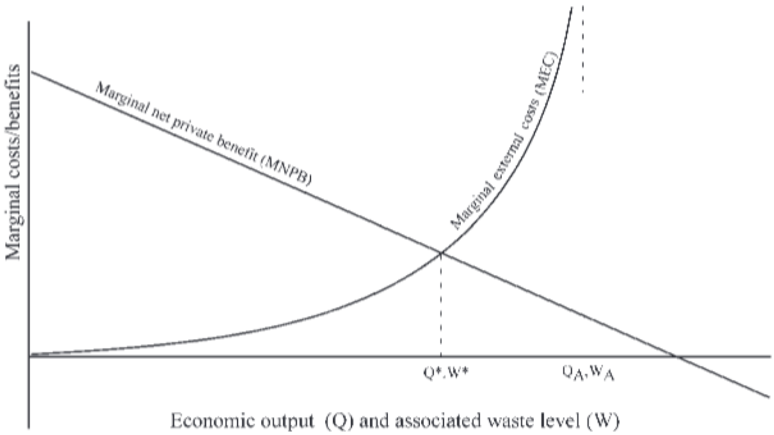

Figure 12.7 • Waste absorption capacity: marginal costs and benefits of pollution.

On the other hand, the extreme levels of uncertainty and ignorance concerning waste absorption capacity and the impacts of paper mill pollutants are not reflected in this graph. In reality, both the MEC curve and the line depicting waste absorption capacity should be thick smears instead of the fine lines depicted. We have labeled the optimal level of output/pollution as Q*,W*, but again caution that with our given state of knowledge, this is a broad range, not a precise point. The reasoning behind the location of Q*,W* was described in Chapter 10.

We cannot stress enough the importance of looking at pollution flows and waste absorption capacities as dynamic. Some economists have argued that pollution causes zero damage before reaching point Q_A, W_A, because the ecosystem is capable of assimilating the waste. But there clearly are substantial costs to pollution even when the ecosystem is capable of assimilating them. MEC may approach infinity at Q_A,W_A.

Depending on the population level and the level of economic activity (i.e., based on scale), MNPB may become zero before or after reaching the waste absorption capacity of the environment. The more full our planet becomes—the larger our scale—the more likely that MNPB will still be positive when we reach the waste absorption capacity and MEC approaches the vertical. Why? Goods and services are characterized by diminishing marginal utility. This means that for a given amount of goods, more people imply lower per-capita consumption, and hence higher utility from each unit consumed. This shifts the MNPB curve upward.

Finally, we must stress once again here that even if policy makers could measure the full marginal costs and benefits of pollution and set the number of permits accordingly, pollution markets would still fail to generate all the wonderful properties associated with the free market. Different individuals obviously have different preferences (utilities) with respect to polluted environments. Markets are widely extolled because they allow the individual to choose what she produces and consumes so that her marginal benefits from either are exactly equal to her marginal costs. Pollution, however, affects public goods, and all individuals must consume the same amount. It would be impossibly complex to create a system in which each individual was paid by the polluter according to his or her own dislike of pollution. This by no means implies we are opposed to markets in waste absorption capacity, but it does mean that we should not associate with them all the market virtues associated with the buying and selling of market goods.

BIOTIC AND ABIOTIC RESOURCES: THE WHOLE SYSTEM

To achieve a sustainable, just, and efficient economic system, we clearly must understand the nature of the resources upon which that system depends. We need to understand the role these resources play in meeting the needs of humans and other species on this planet, and the characteristics that affect their allocation within and between generations via market and nonmarket mechanisms. It would, of course, be impossible to analyze every individual resource. Instead, we introduced the important concepts of rivalness, excludability, externalities, ignorance and uncertainty, and stock-flow and fund-services, and applied these concepts to specific categories of natural resources. To facilitate this, we created a rough taxonomy of biotic and abiotic resources subdivided into eight categories, and applied the above concepts to each of these categories.

Our first goal with this approach was to help you understand precisely why markets fail to efficiently allocate each individual resource and start you thinking about what types of institutions and mechanisms might work better. We began with abiotic resources, which are fairly simple to understand. We then moved on to the stock-flows provided by nature, the raw materials on which the economy depends. We began to see the emergence of complexity: unpredictable ecological thresholds beyond which a population will collapse, impacts from outside variables such as climate change and habitat degradation. The analysis grew a bit more complicated. Once we turned to ecosystem services and waste absorption capacity, it became obvious that these were elements of a whole system and could not be understood apart from that system. Ecosystem fund-services, including waste absorption capacity, are provided by the complex interaction of ecosystem stock-flows, and are necessary to sustain those stock-flows. We can't think of allocating the stock-flow independently of the fund-service. Both fund-service and stock-flow are seriously affected by the waste flows from nonrenewable abiotic resources. So what does this mean?

Our second goal with this approach was, paradoxically, to guide you toward the conclusion that the first goal is insufficient. While it helps to understand the particular characteristics of each individual natural resource, it is more important by far to recognize that these resources are so intimately intertwined that we cannot allocate any one resource without considering the impact it will have on others. The reductionist approach (breaking down low-entropy resources into relatively narrow categories) is important, but inadequate on a complex, living planet. Seemingly efficient allocation of each resource individually will not necessarily lead to the efficient allocation of all resources together. Ecological economics is concerned with integrated systems, not individual commodities, and with complex societies, not atomistic individuals. Breaking a system down to better understand its individual components is a useful analytic tool, but it can seriously mislead us unless we subsequently synthesize these components into an integrated understanding of the whole.

BIG IDEAS to remember

- Stock-flow versus fund-service resources
- Sustainable yield, maximum sustainable yield
- Absorptive capacity
- Per-unit effort curve
- Stable and unstable equilibrium
- Maximizing annual profit versus maximizing net present value
- Exploitation of wild versus bred populations
- Natural dividend