
An Economics Perspective on Treating Voluntary Programs as Clubs

Matthew J. Kotchen and Klaas van 't Veld

The overall theme of this book is to examine how club theory can advance our understanding of voluntary programs, defined as regulatory institutions under which businesses voluntarily agree to comply with environmental or other standards. Drawing on insights from the economics literature on club theory, Potoski and Prakash (2005b; Prakash and Potoski 2006b) develop a political science perspective on how voluntary programs can be interpreted as clubs.¹ Their analyses, along with the contributed chapters in this volume, open a new area of inquiry into the conceptual underpinnings of voluntary programs and the potential for voluntary programs to function as decentralized mechanisms for mitigating collective action problems.

According to one definition, “a club is a voluntary group of individuals who derive mutual benefit from sharing one or more of the following: production costs, the members’ characteristics, or a good characterized by excludable benefits” (Cornes and Sandler [1986] 1996, 347). To see the parallel with voluntary programs, interpret “individuals” to mean firms, and consider the example of forest certification programs that seek to promote sustainable forest practices. The certification program establishes a benchmark of best practices, and firms that satisfy the benchmark, which is assumed to be costly, may seek voluntary certification. The benefits to the firm are reputational, in that consumers trust the certification and may be willing to pay a premium for sustainably harvested forest products. In this instance, the forest certification program is effectively a club that is based on member characteristics (sustainable practices) and offers excludable benefits (reputation). When described in this way, voluntary programs appear to be compatible with much of the existing economic theory on clubs.

Many voluntary programs, however, have an additional feature that as far as we know has not been accounted for in the existing theory. These voluntary programs are designed to promote the spillover of positive externalities outside the club. The intent of forest certification programs, for example, is to promote habitat conservation for external benefits such as the protection of biodiversity and recreational opportunities. Note that these positive externalities—which flow outside the club—are public goods. Hence, in order to fully understand the institutional arrangements of voluntary programs, we need to expand the conceptualization of clubs to account for the broader context of public goods provision. Among the cases discussed in this volume that are consistent with this conceptualization are the Kimberly Process for diamond certification (chapter 5), voluntary labor standards (chapter 6), and government-sponsored environmental programs (chapters 10 and 11). Other examples can be found in Morgenstern and Pizer's (2007) collection of papers on voluntary environmental programs in the United States, Europe, and Japan.

The aim of our chapter is to develop a formal economic model that nests elements of club theory within a model of the private provision of a public good. The model is intentionally simple, yet it provides a useful starting point for addressing a number of fundamental questions raised in chapter 2. Are voluntary clubs an effective mechanism for mitigating collective action problems? How do club standards emerge? How are club standards related to club size? And under what circumstances, if any, will club standards and/or membership be socially optimal?

We motivate the model with a consumption good that can be produced with different levels of a “green characteristic”—namely, the environmental friendliness of practices used in its production. Producers choose the level of this characteristic, whereby greener production is assumed to be more costly. Consumers care about the characteristic, but cannot directly observe it (because, say, they cannot observe production practices). As a result, green production is only credible if it is certified by a “green club” that producers can voluntarily join.

Establishing and managing such a club is itself costly, however, and it is reasonable to assume that the average club costs—the total costs divided by the number of member firms—eventually increase with the club size (for example, due to additional layers of bureaucracy, or greater complexity of monitoring and enforcing the club's green production

standard). This eventual increase in the average club costs plays a key role in our model, because it in effect gives rise to a negative congestion externality. As Sandler and Tschirhart (1997) note, the presence of some form of crowding is a defining premise of club theory, and it has the effect of making the level of whatever benefit the club provides interdependent with the club's membership size. In our model, the congestion externality creates an interdependence between the club standard and the number of firms in the club.

In the next section we present the model, beginning with the special case in which consumers have purely "warm glow" preferences; that is, consumers care only about the private provision of the green characteristic through their own purchase of the green good.² We show that in this case, the socially optimal club standard balances the benefit that consumers of the green good derive against the higher cost of green production, while the socially optimal club size balances the benefit from expanding the club's overall provision of the green good against the higher cost of managing the club. We demonstrate also that in a market setting with no restrictions on club membership, the equilibrium club size will be inefficiently large, to the point where the club makes no net contribution to social welfare. The reason is that when firms decide whether or not to join the club, they consider only their private benefit and cost of doing so, but ignore the congestion externality they impose on existing members.

We then extend the model to account for more general preferences over the green characteristic as a public good, and again compare the socially optimal club with the club that arises in open-access market equilibrium. We show that the positive public good externality of providing the green good increases both the socially optimal club standard and club size. Because firms ignore the positive externality, though, while continuing to ignore the negative congestion externality, the market equilibrium club size may now be either too small, too large, or socially optimal. Nevertheless, the club always contributes to social welfare, because the public good benefit is not dissipated in equilibrium.

Following that, we analyze policies that can improve the equilibrium club's efficiency, but also discuss realistic constraints on such policies. We switch gears to consider an "environmentalist club" that simply seeks to maximize the provision of the public good. Here we show that a government concerned with the different objective of maximizing social

welfare may in some cases find it optimal to simply encourage an environmentalist group to sponsor the club, and then we offer some concluding thoughts.

Warm-up with Warm Glow

Consider an economy of identical producers, each of which is able to produce one unit of a particular consumption good using either conventional production or a green production process with different levels of environmental friendliness. Let $\theta \geq 0$ denote a producer's chosen level of green production, where $\theta = 0$ corresponds to conventional production. Green production is assumed to be more costly, such that the total cost of producing a unit of output is $c + \alpha\theta$, where c is the cost of conventional production and $\alpha\theta$ is the additional cost of green production.

The economy includes N identical consumers. Each consumer is assumed to purchase one unit of output and have preferences of the form

$$U(\theta) = b + f(\theta),$$

where b is the benefit from consuming the conventionally produced good, and $f(\theta)$ is the additional benefit from consuming a good produced at green production level θ . This additional benefit is assumed to be strictly increasing in θ , but at a decreasing rate.³ This form of the utility function is considered a warm-glow specification, because each individual cares only about the green production associated with their own consumption. In the next section, we generalize consumer preferences to account for the public good aspect of green production.

To keep things simple, we assume that green production is not observable to consumers, and that it is not credible for producers to claim green production at any level unless they are certified by a green club. It follows that without a club, no producer would engage in green production (i.e., $\theta = 0$), as they would have no incentive to incur the additional cost. We assume that b exceeds c , so that producing a unit of the conventional good improves social welfare with positive surplus $b - c$. How this surplus is split between consumers and producers depends on the price for the good, which in turn depends on market conditions.⁴ Regardless of the price, however, the level of social welfare—equal to the sum of consumer and producer surplus—would be $N(b - c)$ without a club.

We consider the case of a single green club that requires its members to meet a benchmark standard of θ . Moreover, we assume that the club

can perfectly monitor and enforce this standard, so that there is no issue of shirking by the member firms. The club thus provides a mechanism that makes its members' green production at the benchmark θ fully credible to consumers. But there are costs associated with establishing and managing the club. These include both "fixed" overhead costs, which are incurred regardless of the club size n , and "variable" administration costs, which increase with n . Overall club costs are shared equally among members, so that each must pay the average cost $A(n) = C(n)/n$. Provided, then, that variable costs do not increase too rapidly with n , the fact that fixed costs can be spread over more members will cause the average costs to initially decrease with n . Yet we assume that as n increases further, the increased complexity of administering a large club will cause the average costs to eventually increase with n . As noted in the introduction, this assumption plays a key role in our model. Formally, we assume that there exists a critical \check{n} such that $A'(n) < (>)0$ for $n < (>)\check{n}$.

The Socially Optimal Club

A green club is fully characterized by the combination of its standard θ and membership n . We first consider the $\{\theta, n\}$ combination of the socially optimal club—that is, the combination that would be chosen by a hypothetical social planner concerned with maximizing social welfare. Formally, if we let W denote welfare, the planner's optimization problem would be

$$\max_{\theta, n} W = n[b + f(\theta) - (c + \alpha\theta) - A(n)] + (N - n)(b - c).$$

The first term on the right-hand side represents the overall surplus generated by the n firms in the club that each produce a single unit of the green good. This is equal to the aggregate benefit enjoyed by n consumers, $n[b + f(\theta)]$, less the aggregate cost of production, $n(c + \alpha\theta)$, and the cost of managing the club, $C(n) = nA(n)$. The second term represents the surplus welfare generated by the $N - n$ firms outside the club that each produce a single unit of the conventional good. Rearranging terms, we can rewrite the problem more simply as

$$\max_{\theta, n} W = n[f(\theta) - \alpha\theta - A(n)] + N(b - c), \quad (4.1)$$

where the first term is now the *additional* surplus generated by the club, over and above the surplus $N(b - c)$ that is generated even without the club.

The socially optimal $\{\theta^*, n^*\}$ combination that solves this problem is implicitly defined by the first-order conditions

$$f'(\theta) = \alpha \quad (4.2)$$

and

$$f(\theta) = \alpha\theta + A(n) + nA'(n). \quad (4.3)$$

Equation (4.2) shows that at the optimal combination, the marginal benefit of increasing the standard should equal the marginal cost. Equation (4.3) equates the marginal benefit and marginal cost of increasing the club size. Adding a firm to the club yields a marginal benefit $f(\theta)$, because one additional consumer will consume the green good. At the same time, adding a firm implies a marginal cost $\alpha\theta + A(n) + nA'(n)$, because producing the additional unit of the green good costs $\alpha\theta$, and because the overall club costs increase by $C'(n) = A(n) + n'A(n)$.

Note that as long as the average surplus generated by the club, as given by the term in brackets in expression (4.1), is positive, condition (4.3) will hold only if $nA'(n) > 0$. That is, average club costs must be increasing, implying that the socially optimal club size n^* must be strictly larger than the club size \check{n} at which the average costs are minimized. At any such n , a firm's joining the club increases the overall club costs by *more* than the average club cost $A(n)$ before it joined. The difference, equal to n times the increase $A'(n)$ in all member firms' average costs, is essentially an external cost that the new firm imposes on existing members.

Figure 4.1 illustrates the solution graphically. The figure shows how, at the socially optimal standard θ^* defined by equation (4.2), the average club cost AC varies with club size n , and how this cost compares with the social marginal cost SMC as well as the marginal benefit $f(\theta^*)$ of adding a firm to the club. Beyond club size \check{n} , the average cost increases with the club size, which implies that each additional firm must contribute above-average additional costs—that is, the marginal club cost exceeds the average club cost. The socially optimal club size n^* occurs where the SMC is just offset by the marginal benefit. At that size, the average surplus from the club, $f(\theta^*) - \alpha\theta^* - A(n^*)$, is strictly positive and equal to the external cost $n^*A'(n^*)$ imposed by the last firm that joined.

Equations (4.2) and (4.3) capture how the surplus is optimally created through the exchange of club-certified goods. When green production is

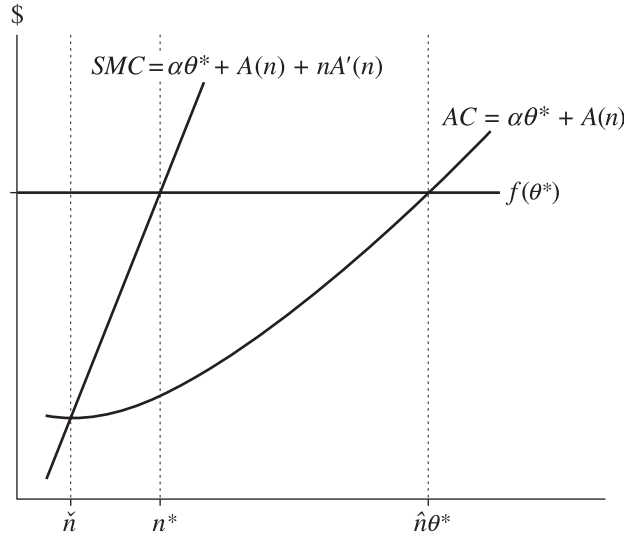


Figure 4.1
Socially optimal and equilibrium club size

credible, consumers are willing to pay a premium up to $f(\theta)$ for each additional unit produced by the club, and this more than covers the increased production costs $\alpha\theta$. Because of this, the optimal standard should maximize the difference between the two, as captured in equation (4.2). Equation (4.3) shows, though, that from the perspective of maximizing social welfare, club membership and thereby the output of the green good should be limited. This is because of two additional costs of expanding club membership and output: the average cost $A(n)$ of club administration and the external cost $nA'(n)$ that each additional member imposes on all other members.

The Equilibrium Club

We now examine the characteristics of the club that will emerge in market equilibrium if the club is open access, allowing any firm that meets its standard to join. In order to analyze this equilibrium, we must consider the price that will emerge for the club-certified good. It turns out that we can focus exclusively on the price premium that will emerge—that is, the difference between the price for the green good and that for the conventionally produced good.

Denoting the premium p , let us take a given club standard θ as a starting point for considering the equilibrium size of the club. Firms have an incentive to join the club as long as the premium they receive exceeds their increase in costs from joining (i.e., $p - \alpha\theta - A(n) \geq 0$). Moreover, consumers have an incentive to purchase the additional firm's output as long as their utility gain from doing so exceeds the price premium (i.e., $f(\theta) - p \geq 0$). It follows that as long as the *overall* surplus from additional green production is positive (i.e., $f(\theta) - \alpha\theta - A(n) \geq 0$), some price premium exists that meets both conditions, allowing a further expansion of the club. Once club membership grows beyond the critical size \check{n} , however, the component $A(n)$ of firm costs will start to increase, until eventually the overall surplus is exhausted and it must hold that

$$f(\theta) = \alpha\theta + A(n). \quad (4.4)$$

No further expansion is then feasible, because it would raise firm costs $\alpha\theta + A(n)$ above the consumer willingness to pay $f(\theta)$, leaving no price premium that can satisfy both sides of the market.

It is important to recognize that equilibrium condition (4.4) does not define a unique set of club characteristics $\{\theta, n\}$. In fact, there are an infinite number of combinations that will satisfy the equation. Nevertheless, if we start with a given club standard θ , then the mapping to an equilibrium club size n is unique. We can write this mapping as a function $\hat{n}(\theta)$, which is illustrated in figure 4.2. Notice that no equilibrium club exists if the standard is set lower than $\underline{\theta}$ or higher than $\bar{\theta}$. This follows because at both low and high standards, the difference between consumers' willingness to pay and green production costs is too small to cover the average club costs, even at the club size \check{n} where the average club costs are minimized.

Figure 4.2 also shows how $\hat{n}(\theta)$ is inversely U shaped. The intuition for this result is straightforward. For low-standard clubs, raising the standard raises consumers' willingness to pay for green production at a faster rate than the cost of green production. It follows that more firms have an incentive to join, even though the average club costs are increasing. In contrast, for high-standard clubs, consumers' willingness to pay increases at a slower rate than the cost of green production, so firms must drop out in order to reduce the average club costs.

Formally, it can be shown that the slope of $\hat{n}(\theta)$, or the change in the equilibrium club size for any change in the club standard, is




Figure 4.2
Equilibrium club size as a function of the club standard

$$\hat{n}'(\theta) = \frac{f'(\theta) - \alpha}{A'(\hat{n})}. \tag{4.5}$$

Comparing the numerator of this expression with equation (4.2) implies that the slope of the function equals zero at θ^* . Hence, equilibrium club membership is maximized at the socially optimal club standard. Note that the equilibrium size $\hat{n}(\theta^*)$ at this standard is strictly greater than the socially optimal club size n^* indicated by the point W^{\max} . It follows that the socially optimal club cannot arise in open-access market equilibrium.

Recall from figure 4.1 that at the efficient club size n^* , the consumers' benefit from an additional unit of the green good equals the social marginal cost of having a new club member. In the open-access equilibrium, however, potential club members consider only their *private* marginal cost (equal to the average cost AC), which excludes the external cost that their joining imposes on other members. They thus continue to join as long as the consumers' benefit, or willingness to pay, exceeds the average cost, which is true up to club size $\hat{n}(\theta^*)$.

At this higher, equilibrium club size, the surplus generated by the club is completely dissipated. That is, the equilibrium club makes no contribution to social welfare. Moreover, this is true not just at the socially optimal standard but also at any standard that the club might set. This

constitutes a market failure in economic terms. Rather than maximizing overall surplus, leaving the market to determine the club size ends up driving the surplus to zero. The underlying mechanism is essentially that of the well-known “tragedy of the commons.” Just as herders ignore the external costs that their grazing imposes on fellow herders, resulting in overgrazing, firms ignore the external costs that their joining the club imposes on other members, resulting in excess entry.

Generalization to Public Good Preferences

We discussed previously how many voluntary programs seek to promote the provision of public goods. Thus far we have assumed only warm-glow benefits from voluntary provision. This section expands the model to account for public good benefits that might arise from green production.

The model setup remains unchanged, except for the specification of consumer preferences. To account for both warm-glow and public good benefits, we now assume that consumers have preferences of the form

$$U(\theta) = b + f(\theta) + g(n\theta),$$

where the term $g(\cdot)$ increases at a decreasing rate in the overall production of the green good $n\theta$. This term captures public good benefits that result from overall green production, in the form of nonrival and non-excludable improvements to environmental quality that all N consumers enjoy equally, regardless of whether they contribute to such production by purchasing the good.⁵ We assume additive separability between warm glow and the public good for analytic simplicity.

The Socially Optimal Club

We analyze the public good version of the model following the same steps as those above. The social planner’s objective is to solve the following maximization problem:

$$\max_{\theta, n} W = n[f(\theta) - \alpha\theta - A(n)] + N[g(n\theta) + b - c]. \quad (4.6)$$

The first-order conditions, which again uniquely define the socially optimal club characteristics $\{\theta^*, n^*\}$, are now

$$f'(\theta) + Ng'(n\theta) = \alpha \quad (4.7)$$

and

$$f(\theta) + Ng'(n\theta)\theta = \alpha\theta + A(n) + nA'(n). \quad (4.8)$$

Comparing equations (4.2) and (4.7) shows that the social marginal benefit from increasing the club standard now includes an additional term, $Ng'(n\theta)$, which captures the social marginal benefit to all N consumers of increasing the public good. Similarly, comparing equations (4.3) and (4.8) reveals that the marginal benefit of increasing the club size now includes an additional term, $\theta Ng'(n\theta)$, which captures the social marginal benefit to all N consumers from having one additional firm produce the green good.

Using equations (4.7) and (4.8), it is straightforward to verify that both θ^* and n^* are strictly greater than in the pure warm-glow case of the previous section. Driving this result are the public good benefits that all consumers now enjoy from the club, regardless of whether they purchase the club-certified good. In effect, the club is now more socially beneficial because it generates public good spillovers to all individuals in the economy.

The Equilibrium Club

We can show that the change in consumer preferences to account for public good benefits has no effect on the forces that drive the equilibrium club size. Firms still have an incentive to join the club as long as the price premium p exceeds their private costs $\alpha\theta + A(n)$. Moreover, consumers still have an incentive to purchase the green good as long as their private warm-glow benefit from the purchase, $f(\theta)$, exceeds the price premium p . The reason why the public good benefit plays no role in this decision is that for any reasonable club size n , consumers receive the same level of the public good $n\theta$ regardless of whether they themselves contribute. This is captured implicitly with our assumption that n is continuous. It follows that the equilibrium club size continues to be determined by equilibrium condition (4.4).

Comparing the equilibrium condition (4.4) with the social planner's first-order condition (4.8), we now see two sources of market failure. The first, which we described as a tragedy of the commons, is unchanged from the warm-glow version of the model. When firms decide whether to join the club, they still ignore the negative externality that they impose on other club members. The new source of market failure arises because of free riding with respect to the private provision of the public good. Because consumers enjoy the public good regardless of whether they

themselves purchase the club-certified good, they have no additional willingness to pay for the good beyond that derived from warm glow. Accordingly, when firms continue to join the club, they ignore the positive externality, $Ng'(n\theta)\theta$, that they impose on all consumers.

In the special case where the two externalities are exactly equal at $\{\theta^*, n^*\}$, the two market failures will offset each other, and the equilibrium club size for standard θ^* will equal n^* . More generally, however, the two externalities will differ in size, resulting in an equilibrium club that is suboptimally large or small.

Finally, it is important to recognize that the public good case differs from the pure warm-glow case because the equilibrium club does in fact contribute to social welfare. The reason is that even though the warm-glow surplus $n[f(\theta) - \alpha\theta - A(n)]$ is still driven to zero through excess firm entry, all consumers still enjoy positive public good benefits $Ng(n\theta)$. In effect, the warm-glow benefit to the n consumers that purchase the club-certified good “pays” for the public good benefit to all N consumers. Consumer purchases of warm glow, in other words, produce a positive externality.

Club Policies

Up to this point, we have considered how a hypothetical social planner would optimally choose both the club standard and size, and how given any club standard, market forces will determine the equilibrium club size. We have not yet considered how a real-world administrator—possibly a government agency, an NGO, or some other third party—would choose the club standard in anticipation of the equilibrium response. Nor have we considered how the administrator might take advantage of various policy instruments. In order to analyze the latter questions, we must specify the administrator’s objective function. In this section we assume that the administrator aims to maximize social welfare, just like the social planner. In the next section we consider an alternative.

We have already shown that in order to implement the welfare-maximizing club combination $\{\theta^*, n^*\}$, it is not in general sufficient to simply set the standard at θ^* . Without any restrictions on club entry, a tragedy of the commons will tend to increase the club size above n^* , while free riding by consumers will tend to reduce the club size below

n^* . That being said, in many circumstances it may be difficult legally or politically to impose restrictions on a club's size.

A more practical alternative may be to charge a uniform admission fee τ , over and above the average club administration costs $A(n)$. Since this fee just acts as an additional cost for firms, which all else equal reduces their incentive to join, the equilibrium condition (4.4) becomes

$$f(\theta) = \alpha\theta + A(n) + \tau. \quad (4.9)$$

Comparing this condition with the social planner's first-order condition (4.8) shows that setting the fee at level

$$\tau^* = n^* A'(n^*) - Ng'(n^* \theta^*) \theta^* \quad (4.10)$$

will implement the socially optimal club. In effect, the admission fee forces entering firms to internalize the two externalities associated with membership. The first term on the right-hand side of equation (4.10) captures the negative externality of increasing average club costs for all other club members. The second term captures the positive externality of increasing the level of the public good for all consumers. Depending on which of these terms is larger, τ^* may be either positive or negative. In the latter case, the optimal policy would be to subsidize club membership, thereby raising it to n^* from what would otherwise be suboptimally low.

An equivalent mechanism to the admission fee is a tax τ on the green good. Such a tax drives a wedge between the premium p_c paid by consumers and the premium p_f received by firms, such that $p_c = p_f + \tau$. Nevertheless, club entry will as always continue until both $f(\theta) - p_c = 0$ and $p_f - \alpha\theta - A(n) = 0$. Combining these three conditions and substituting away p_c and p_f leaves exactly the same equilibrium condition (4.9). Hence, setting tax τ as in equation (4.10) would also implement the socially optimal club, and again τ^* can be negative, implying a subsidy on the green good.

An important concern relating to the different policy instruments has to do with revenue. Specifically, if $\tau^* > 0$, the aggregate revenue raised from the club through either an admissions fee increase or a tax on the club good cannot be returned to either the n member firms or the n consumers buying from the club without canceling the policies' incentive effects. The revenue must instead be used in a manner that benefits *all* consumers or firms equally, regardless of whether they respectively buy

or produce the green good. Similarly, if $\tau^* < 0$, the aggregate revenue $n\tau^*$ required to finance the admissions fee reduction or subsidy on the club good must be raised from all consumers or firms equally. Alternatively, revenue could be raised in some other sector of the economy, but then one would need to consider potential inefficiencies there.

These constraints may in practice make club policies less attractive or even infeasible. If the whole rationale for establishing a “voluntary” club is that it provides an alternative to government regulation, whether through command-and-control mandates, taxes, fees, or subsidies, then the program may have to operate under what amounts to a budget-balancing constraint: any fees associated with club membership must be set at levels that neither exceed nor fall short of covering club administration costs. In such contexts, if the club administrator *can* choose the club size directly, it will be constrained from choosing any size greater than that consistent with the equilibrium condition (4.4). In other words, given whatever standard θ the administrator chooses, it will be constrained to set n less than or equal to the club size $\hat{n}(\theta)$ implicitly defined by condition (4.4).

If the club administrator *cannot* choose the club size directly at all, then a budget-balancing constraint will reduce its role to simply choosing the club standard θ , knowing that the resulting club will be of the equilibrium size $\hat{n}(\theta)$. This implies that in the pure warm-glow case discussed earlier in this chapter, a welfare-maximizing club administrator will be indifferent between all standards within the interval $[\underline{\theta}, \bar{\theta}]$, and what is more, between establishing a club and not establishing one. This is because, as shown previously, the equilibrium club for any θ will make no contribution to social welfare. In contrast, in the public good case examined earlier, we found that the equilibrium club does contribute to social welfare. Even though club-related benefits are driven to zero, the club generates public good spillovers $Ng(n\theta)$ to all individuals in the economy. The implication is that the welfare-maximizing club administrator will choose θ to maximize the public good spillovers.

Environmentalist Clubs

Having considered a club whose administrator aims to maximize social welfare, we now explore the alternative of a club whose administrator aims to simply maximize the provision of the public good. Analyzing this objective seems appropriate for a variety of scenarios, including

ones in which the public good is an environmental amenity and the administrator is not a government agency but an environmental group.

In terms of the model presented here, an environmental group may establish a club that certifies a particular good as green, with the explicit objective of maximizing $n\theta$. In principle, the group may be able to choose not just the club characteristics $\{\theta, n\}$ but also the premium p charged for the club-certified good. Following the same reasoning used above, we find that the familiar equilibrium condition must continue to hold. We can thus write the environmentalist club's problem as follows:

$$\max_{\theta, n} n\theta \quad \text{subject to} \quad f(\theta) = \alpha\theta + A(n)$$

In other words, the club seeks to maximize the provision of the public good subject to the constraint of the familiar equilibrium condition.

Figure 4.3 illustrates the problem graphically. The downward-sloping curves labeled $G_0, G_1,$ and so on, represent the club's indifference curves—that is, combinations of n and θ that yield a particular value $G_i = n\theta$ of its objective function, with higher curves yielding higher values ($G_0 < G_1 < G_2 < G_3$). As in figure 4.2, the bold, inverse U-shaped curve represents the function $\hat{n}(\theta)$ implicitly defined by the equilibrium




Figure 4.3
The environmentalist club's optimization problem and socially optimal combinations with public good preferences

condition. Graphically, the club's problem involves choosing a standard θ such that the associated point on the constraint $\hat{n}(\theta)$ reaches its highest indifference curve, thereby maximizing $n\theta$ subject to the equilibrium condition. This is achieved at the point $\{\theta^e, n^e\}$, where the constraint is tangent to indifference curve G_1 .

Although figure 4.3 is drawn for a particular parameterization of the model, we can describe some general results comparing the environmentalist club with the socially optimal club. We know that because the club's indifference curves are downward sloping, the tangency point that defines $\{\theta^e, n^e\}$ must lie on the downward-sloping segment of the $\hat{n}(\theta)$ curve. This implies that the tangency must lie to the right of the socially optimal standard when preferences are purely warm glow, because we saw earlier (recall figure 4.2) that the latter standard is where the $\hat{n}(\theta)$ curve is maximized. In the purely warm-glow case, in other words, the environmentalist club will choose a standard θ^e higher than the welfare-maximizing standard θ^* . It can also be shown that the environmentalist club size n^e will be greater than the welfare-maximizing size n^* .

When preferences include a public good component, these results continue to hold, but only if N is not too large. Recall from the analysis above that both θ^* and n^* increase from their values in the pure warm-glow case. This increase can be shown to be larger the greater that N is. It can also be shown that at a critical value of N , which we denote N^e , the socially optimal $\{\theta^*, n^*\}$ coincides with the environmentalist club's optimal $\{\theta^e, n^e\}$. For values of N greater than N^e , the relationship between the two optima becomes reversed—that is, the environmentalist club standard and size both become smaller than would be welfare maximizing. The reason for the difference is that the environmentalist club is constrained by the equilibrium condition, while the social planner is not.

The upward-sloping curve in figure 4.3 from $W^{\max}(0)$ to $W^{\max}(N)$ illustrates these results graphically. The curve is the locus of socially optimal combinations of $\theta^*(N)$ and $n^*(N)$ when these are treated as functions of N . The lowest point on the curve, labeled $W^{\max}(0)$, corresponds to the socially optimal club in the warm-glow case, when the public good benefits $Ng(n\theta)$ are zero. When the public good benefits are positive, the social optimum will lie further up the curve, and more so the larger N is. Clearly, then, there exists a critical N^e such that the social optimum will coincide with point $\{\theta^e, n^e\}$, and at higher N it will lie above and to the right of that point.

An important observation about the environmentalist club relates back to the discussion in the previous section about a real-world club administrator whose aim is to maximize welfare. It was noted there that if such an administrator faces a budget-balancing constraint, and in particular cannot subsidize the club with outside funds, it will not be able to choose a club size that exceeds the equilibrium size $\hat{n}(\theta)$ conditional on whatever standard it sets. In terms of figure 4.3, the administrator's choices would be limited to the segment of the W^{\max} locus below the $\hat{n}(\theta)$ curve. It follows that for N greater than N^e , when the administrator's unconstrained optimum would lie above the $\hat{n}(\theta)$ curve, its constrained optimum will coincide with the environmentalist club outcome. Interestingly, this provides the administrator with an alternative policy option: rather than creating and administering the club itself, it can simply encourage an environmental group to create the club, and achieve the same constrained-optimal outcome.

The same implication follows when the welfare-maximizing administrator cannot choose the club size directly, in addition to facing a budget-balancing constraint. As also discussed in the previous section, the administrator's effective optimization problem then reduces to maximizing public good spillovers alone, subject to the equilibrium condition. But since those public good spillovers $Ng(n\theta)$ increase monotonically in $n\theta$, that optimization problem is effectively identical to that of the environmentalist club, with the same solution.

Conclusion

The economic model developed in this chapter offers a starting point for thinking formally about voluntary programs as clubs, nested within the context of public goods provision. Voluntary programs are treated as clubs because they provide nonrival but excludable benefits to members, yet these benefits are subject to negative congestion externalities when the club membership expands. Because the club monitors and certifies certain production practices that consumers value but cannot themselves verify, they allow member firms to earn a premium for their goods. But these profits are reduced by shared club administration costs, which eventually increase as the club membership expands. The model is distinct from other applications of club theory in that the club promotes the spillover of positive externalities. In the context of clubs that promote

green production practices, this positive externality consists of an environmental public good.

In the special case where consumers derive only private, warm-glow utility from the club good, there are no public good spillovers. If club membership is open to all firms that agree to meet its certified production standard, the congestion externality internal to the club gives rise to a tragedy of the commons: new firms will continue to join the club as long as there are profits to be made. The result is a market failure due to the complete dissipation of the potential benefits of the club.

In the more general case where consumers derive both private and public good utility from the club good, the same tragedy of the commons can arise. But in this case, there is a further source of market failure because of free riding among consumers. Consumers have no incentive to consider the public good benefits to others of their own purchases of the club good. We find that the first market failure still results in an equilibrium where direct club benefits are completely dissipated. The second market failure, however, still leaves an indirect benefit of the club: consumers' purchases of the club good for warm-glow reasons end up "paying" for a public good benefit to all consumers—even those who do not purchase the club good. Moreover, because the second market failure tends to discourage firm entry relative to what is socially optimal, the equilibrium club size with both market failures combined may be either greater than, less than, or equal to the welfare-maximizing size.

We also considered several policies that welfare-maximizing club administrators might employ in order to address the two market failures. Most obviously, these include direct limits on club size, or if such limits are infeasible, taxes or subsidies that discourage or encourage firm entry. If the latter policies are infeasible as well, perhaps because they are viewed as inconsistent with the "voluntary" nature of the club, we find that the club administrator's constrained optimal policy reduces to simply maximizing the overall provision of the public good. Interestingly, then, its optimal policy may be to leave club administration to an outside group whose objective is to maximize such provision to begin with, rather than to maximize welfare. Third-party certification and ecolabeling programs are examples.

In conclusion, we hope that the model presented here provides a useful start for developing formal models that capture the institutional arrangements of voluntary programs. Although the setup of our model is inten-

tionally simple, the treatment of club formation within the context of the private provision of a public good generates several new insights. Further extensions not considered here, but left for future work, include the consideration of imperfect monitoring and the enforcement of club standards, heterogeneity among consumers and firms, rival clubs that may arise, cases in which firms can credibly signal their own standards, and game-theoretic strategies.

