
ARTICLE

NONRIVALRY AND PRICE DISCRIMINATION IN
COPYRIGHT ECONOMICS

JOHN P. CONLEY[†] & CHRISTOPHER S. YOO^{††}

INTRODUCTION	101
I. A PRIMER ON PUBLIC GOOD ECONOMICS	105
A. <i>The Critique of Nonexcludability</i>	106
B. <i>The Reconceptualization of Nonrivalry</i>	107
II. A FORMAL MODEL OF PURE PUBLIC GOODS.....	111
A. <i>Perfect Price Discrimination</i>	114
B. <i>Imperfect Price Discrimination</i>	118
III. SPATIAL COMPETITION/IMPURE PUBLIC GOODS.....	120
CONCLUSION.....	128

INTRODUCTION

Legal scholarship on the economics of copyright has largely settled into a debate between two polar extremes. On one side are copyright “neoclassicists,” who favor the expansion of copyright protection until it encompasses all of the present and future uses associated with a creative work, as well as reforms that facilitate price discrimination, on the grounds that innovation is best promoted if authors are able to appropriate as much of the value of their creation as possible. On the

[†] Professor of Economics, Vanderbilt University.

^{††} Professor of Law and Communication and Founding Director of the Center for Technology, Innovation, and Competition, University of Pennsylvania. We would like to thank Joel Waldfoegel, who served as the commentator on our paper, and the other participants in the Symposium on the Foundations on Intellectual Property Reform held at the University of Pennsylvania Law School on January 16-17, 2009.

other side are copyright “minimalists,” who favor limiting the number of uses contained within a copyright so that it provides only enough incentive for innovation and who are generally hostile toward reforms that facilitate price discrimination.¹

Despite the differences in their conclusions, both sides generally frame the arguments in largely economic terms.² Indeed, both sides of the debate analyze copyright through the lens of public goods theory, which Paul Samuelson was among the first to analyze with mathematical rigor.³ A core policy implication of public goods theory is that markets tend to produce too few public goods and underutilize those that are produced.⁴ In the context of copyright, the economic analysis has focused almost entirely on the premise that not only are creative works nonrival in general, but also that any number of additional copies can be produced at zero marginal cost. In so doing, the current literature fails to capture the key economic features that give public goods their distinctive characteristics.

Samuelson imagined that there existed an exogenously fixed set of public goods and pointed out that markets would not produce public goods efficiently even if they were fully excludable.⁵ He derived an optimality condition, now known as the “Samuelson condition,” that requires that the level of a public good be chosen so that the sum of the marginal benefits derived by everyone who consumes the public good equals the marginal cost of production. Thus, satisfying the Samuelson condition in a market context requires determining the marginal benefit that each individual agent obtains from consuming

¹ See Neil Weinstock Netanel, *Copyright and a Democratic Civil Society*, 106 YALE L.J. 283, 311-24, 336-41 (1996) (dividing copyright scholarship into neoclassicist and minimalist schools and providing an overview of their positions).

² See *id.* at 287 (noting that “[s]ome minimalist critics follow the same criterion of allocative efficiency as the neoclassicists”); see also James Boyle, *Enclosing the Genome: What Squabbles over Genetic Patents Could Teach Us*, in PERSPECTIVES ON PROPERTIES OF THE HUMAN GENOME PROJECT 97, 119 (F. Scott Kieff ed., 2003) (noting that intellectual property minimalists are “those most committed to the claim that only utilitarian arguments about the encouragement of future innovation are legitimate parts of the discourse” and exhibit a “hard-wired reflex” to respond to noneconomic arguments by “restrict[ing] their analysis to the . . . utilitarian domain”).

³ For the seminal work, see Paul A. Samuelson, *The Pure Theory of Public Expenditure*, 36 REV. ECON. & STAT. 387 (1954).

⁴ See Christopher S. Yoo, *Copyright and Public Good Economics: A Misunderstood Relation*, 155 U. PA. L. REV. 635, 637 & n.1 (2007) (collecting commentary applying public goods theory to copyright in this manner).

⁵ See Paul A. Samuelson, *Aspects of Public Expenditure Theories*, 40 REV. ECON. & STAT. 332, 335-36 (1958).

the last increment of a public good. Since agents are charged in proportion to these benefits, they have an incentive to hide or underreport these benefits. Although economists have proposed a wide range of potential solutions to this problem, they have largely failed to produce a practical, incentive-compatible mechanism to induce consumers to reveal the intensity of their preferences.⁶

Reconceiving the application of public goods theory to copyright in this manner provides an important middle ground between the extreme positions that dominate the current debate. On the one hand, by underscoring that some degree of price discrimination is essential for the optimal provision of public goods, a more fundamental understanding of public good economics reveals that copyright minimalists' reflexive hostility toward price discrimination is misplaced. Indeed, it is quite likely that prohibiting or inhibiting authors' ability to price-discriminate will reduce economic welfare.

On the other hand, the more refined conception of public goods theory that we propose in this context challenges the neoclassicists' claim that efficient production of public goods depends on authors' ability to appropriate the entirety of the surplus created by their works. Instead, the Samuelson condition implies that when public goods are divisible (in that producers can vary the quantity of public good they produce), producers need only appropriate the *marginal* benefit (rather than the total benefit) resulting from further increases in the production of public goods. This means that some proportion of the available surplus remains with the consumer.⁷ Our approach thus strikes a middle ground between the polar positions that dominate the debate and provides a basis for analyzing the optimal level of price discrimination.⁸

There are two additional complicating factors that make copyright goods a particularly difficult problem even within the context of public goods theory. First, the types of public goods that Samuelson had

⁶ See generally Jean-Jacques Laffont, *Incentives and the Allocation of Public Goods*, in 2 HANDBOOK OF PUBLIC ECONOMICS 537 (Alan J. Auerbach & Martin Feldstein eds., 1987) (surveying and critiquing the literature proposing mechanisms for inducing an efficient production and allocation of goods).

⁷ See Yoo, *supra* note 4, at 674-75 (noting that the Samuelson condition does not require producers to "capture all of the consumer surplus" but rather that producers only capture "the marginal rate at which each consumer would substitute further expansion of the public good for other goods").

⁸ Note that when producers cannot vary the amount of the public good they produce, optimality may require that they appropriate the entirety of the available surplus. See *infra* Part III.

in mind when he developed his theory—such as roads and national defense—are divisible, in that one can provide one unit, two units, or any amount of these goods one chooses. Copyright goods, however, tend to be indivisible. That is, one can either write a book, compose a song, make a movie, or not do so. One cannot write half a book or two units of the same book.⁹ Half a book (perhaps a book written with half the effort) would be a different copyright good, not half the quantity of the same good. Thus, copyright goods involve a fixed first-copy cost and then close-to-zero cost in allowing additional agents to enjoy subsequent copies, in contrast to the zero fixed costs and positive marginal cost typical of classical public goods.

Second, Samuelson followed classical general-equilibrium theory in assuming that the set of public goods is exogenously fixed. However, the major argument for providing copyright protection in the first place is that protection is necessary to provide incentives to produce new copyright goods. Thus, the set of copyright goods is endogenously determined in equilibrium. Economics has only a primitive understanding of how markets create new goods, but this is an issue that cannot be ignored in the context of copyright.

The interplay of these two factors greatly complicates the problem and makes it difficult to make definitive policy recommendations. Indeed, one reason for the persistence of the debate in law is that there is truth in both sides of the argument. Copyrighted works are often imperfect substitutes for one another. This creates demand interactions that change the analysis in important ways. Imperfect substitution adds a different dimension along which public goods can be underprovided: not only can there be too little of a particular good, but markets may also produce too few types of public goods. Perhaps even more surprising is the fact that if policy instruments are not correctly calibrated, it is also entirely possible that markets will provide too many works instead of too few.¹⁰ Thus, conceiving of the copyright markets as endogenously determining the set of discrete public

⁹ We are thus using the term “indivisible” to refer to an aspect of public goods *supply*, in contrast to Kenneth Arrow and other leading public goods theorists who use the term to refer to an aspect of public goods *demand*. See Kenneth J. Arrow, *Economic Welfare and the Allocation of Resources for Invention*, in NAT’L BUREAU OF ECON. RESEARCH, *THE RATE AND DIRECTION OF INVENTIVE ACTIVITY: ECONOMIC AND SOCIAL FACTORS* 609, 615-16 (1962); see also RICHARD CORNES & TODD SANDLER, *THE THEORY OF EXTERNALITIES, PUBLIC GOODS, AND CLUB GOODS* 8 (2d ed. 1996).

¹⁰ For a preliminary analysis of this possibility, see Christopher S. Yoo, *Copyright and Product Differentiation*, 79 N.Y.U. L. REV. 212, 256-64 (2004) (discussing the potential problem of “excessive entry caused by demand diversion”).

goods to provide through a process of monopolistic competition undercuts the seemingly intuitive policy inference that markets invariably tend to produce too few works. To the extent that the relevant equilibrium is one in which overproduction occurs, this insight in turn suggests that the tradeoff between incentives for efficient creation and efficient access to those works that are created—the tradeoff that frames most of the scholarship on the economics of copyright—may be fundamentally misplaced.

The purpose of this Article is to help disentangle these effects and clearly highlight the factors that should guide this policy debate.¹¹ The balance of the discussion is organized as follows: Part I briefly reviews the basic economics of pure public goods and summarizes the shortcomings of the way the debate is currently framed.¹² Part II lays out a basic model of pure public goods and analyzes the role that price discrimination plays in maximizing economic welfare. Part III extends the model by allowing for the possibility of competition between similar products that are imperfect substitutes for one another.

I. A PRIMER ON PUBLIC GOOD ECONOMICS

Under the classic definition, pure public goods have two defining characteristics.¹³ First, they are nonexcludable, which means that the good cannot be provided to one consumer without simultaneously providing it to others. That is, it is infeasible to prevent any agent from consuming whatever public goods are produced. Second, they are nonrival, which means that the consumption of the good by one consumer does not reduce the supply available for consumption by others.

One oft-cited example of a public good is a lighthouse. Lighthouses are said to be nonexcludable, in that it is generally thought that lighthouse services cannot be provided to one ship without also providing them to others. The inability to internalize these positive externalities is said to lead to systematic underproduction of lighthouse services. Lighthouses are also said to be nonrival in that con-

¹¹ This Article represents a formalization of ideas first discussed in an earlier article. *See generally* Yoo, *supra* note 4.

¹² Part I recaps the descriptive account of public good economics presented in the article referred to above. *See id.* at 643-60, 662-75. Those familiar with the subject may wish to proceed directly to Part II.

¹³ *Id.* at 637; *see also* Christopher S. Yoo, *Rethinking the Commitment to Free, Local Television*, 52 EMORY L.J. 1579, 1594 & n.21 (2003) (collecting sources on the origin of the classic definition of pure public goods).

sumption of lighthouse services by one ship does not preclude other ships from consuming them as well. Another example is a fireworks display. It is difficult to allow one consumer to enjoy a fireworks display without simultaneously allowing others to do so. In addition, the enjoyment of the fireworks display by one person does not generally reduce the ability of others to enjoy it as well.

As noted earlier, creative works are generally thought to be pure public goods. Absent some form of copyright protection, authors would be unable to prevent one purchaser from copying the work and sharing it with others. Moreover, the sharing of the work with one person does not in any way reduce the supply of the work available for sharing with others.

A. *The Critique of Nonexcludability*

Interestingly, Samuelson did not regard nonexcludability as an essential characteristic of pure public goods.¹⁴ A moment's reflection on the classic examples discussed above reveals why. Consider the lighthouse. Metering access to a lighthouse is not impossible; it is simply very costly. For example, a lighthouse owner could engage a flotilla of ships to intercept other ships and ensure that only those that had paid for the lighthouse services were able to enjoy them. Alternatively, the lighthouse could install a lamp that operated on frequencies outside the visible spectrum and give only those ships that had paid for the services the equipment necessary to detect its signals. In the context of copyright, the advent of digital-rights management is increasingly allowing authors to charge for uses that were previously nonexcludable.

Stated somewhat more broadly, the problems of nonexcludability are technological rather than fundamental. They are contingent on the underlying technology and may be ameliorated as the underlying technological context changes. In addition, as Coase pointed out, the fact that port usage constitutes a useful proxy for the consumption of lighthouse services allowed lighthouse owners in early England to avoid the problems of nonexcludability simply by including charges for lighthouse services in port fees.¹⁵

¹⁴ See Samuelson, *supra* note 5, at 335 ("Being able to limit a public good's consumption does not make it a true-blue private good.")

¹⁵ See R.H. Coase, *The Lighthouse in Economics*, 17 J.L. & ECON. 357, 364, 375 (1974) (describing how tolls levied on ships benefiting from lighthouses were used to finance private provision of lighthouse services). For an overview and an application to Inter-

B. The Reconceptualization of Nonrivalry

Similar questions have been raised about the way the literature has framed nonrivalry, which, as noted earlier, occurs when consumption by one person does not diminish the supply available for consumption by others. The existing literature models nonrivalry in copyright goods by assuming that once authors incur fixed costs needed to make the first copy of a creative work, the work can be costlessly reproduced an infinite number of times. In other words, the marginal cost of making copies of any existing copyrighted work is zero.

The assumption of zero marginal costs in turn gives rise to the familiar pricing problem that drives much of copyright analysis. Economic analysis suggests that an additional copy of a work should be created whenever the social benefits exceed the social costs of doing so. This occurs when price is set equal to marginal cost. When marginal cost is zero, however, efficient pricing requires that price also be set at zero. Of course, this generates no revenue and prevents producers from recovering the fixed costs of producing the first copy of the work. On the other hand, allowing authors to recover fixed costs by charging prices that exceed marginal cost creates deadweight loss by denying access to potential consumers even though the value they would derive from consuming the work would exceed the cost of allowing them to do so.

Thus, the literature generally views copyright as involving a trade-off between providing sufficient incentives and allowing efficient levels of access to the work that is necessarily second best in both dimensions. Any price that allows authors to recover first-copy costs necessarily reduces access below welfare-optimizing levels. Framing the issue in this manner has caused some to suggest that market failure is so endemic as to be useless as a guide in determining the scope of copyright protection.¹⁶ Other scholars have attempted to solve the marginal-cost pricing problem by proposing “prize” systems under

net pricing, see Christopher S. Yoo, *Network Neutrality and the Economics of Congestion*, 94 GEO. L.J. 1847, 1873-74 (2006).

¹⁶ See Glynn S. Lunney, Jr., *Fair Use and Market Failure: Sony Revisited*, 82 B.U. L. REV. 975, 996 (2002) (arguing that the fact that copyrighted works' public goods character makes market failure inevitable renders market failure a poor guide for determining which uses constitute fair use).

which the government subsidizes the first-copy costs and the resulting works are then sold at marginal cost.¹⁷

As Samuelson pointed out, institutional mechanisms that permit public goods to be priced at marginal cost are not sufficient to solve the problem.¹⁸ Instead, Samuelson focused on another feature: the fact that each person who purchases the public good simultaneously consumes the entire output of the public good.

When goods are rival, the efficient solution is for the price of the good to equal both the marginal benefit each consumer derives from consuming that good as opposed to some other good and the marginal cost of producing that good. In other words, economic welfare is maximized at the point where each consumer faces the same marginal rate of substitution between that good and all other goods and that marginal rate of substitution equals the marginal rate of transformation between that good and all other goods (i.e., $MRS_{\lambda} = MRS_{\lambda} = \dots = MRS_i = MRT$ for consumers $1, \dots, i$).

Markets can provide these goods efficiently because consumers do not have any incentive to misrepresent the intensity of their preferences. Every consumer pays the *same price* and signals the intensity of their preferences by purchasing *different quantities* of the good. Once the market sets the uniform price, the only way consumers can attempt to understate the intensity of their preferences is by purchasing less than their optimal quantity. Consumers have no incentive to do so, since this would only reduce the utility that they would enjoy.

The situation changes when every purchaser consumes the entire industry output. Production increases not by producing additional units, but rather by increasing the size of the public good. For example, a lighthouse may be made taller or may be outfitted with a

¹⁷ For the classic argument in favor of government subsidies of high fixed-cost, low marginal-cost goods, see Harold Hotelling, *The General Welfare in Relation to Problems of Taxation and of Railway and Utility Rates*, 6 *ECONOMETRICA* 242 (1938). The literature on prizes is much more fully developed with respect to patent law. See, e.g., Michael Abramowicz, *Perfecting Patent Prizes*, 56 *VAND. L. REV.* 115, 127-70 (2003) (surveying the literature on patent prizes and proposing another variant). For a similar proposal with respect to copyright, see Steve P. Calandrillo, *An Economic Analysis of Property Rights in Information: Justifications and Problems of Exclusive Rights, Incentives To Generate Information, and the Alternative of a Government-Run Reward System*, 9 *FORDHAM INTELL. PROP. MEDIA & ENT. L.J.* 301, 334-44 (1998).

¹⁸ See Samuelson, *supra* note 5, at 336 (“It is not enough in the decreasing cost case to come closer to marginal cost pricing in the Lerner-Lange manner, making up the deficits by general taxation. As soon as decreasing cost and diversity of product appear, we have the difficult non-local ‘total conditions’ to determine what finite mix of product is optimal.”).

stronger lamp. The number of rockets in a fireworks display can be increased. To use an example based in copyright, the quality of a movie may be increased by hiring better talent, incorporating better screenwriting and music, and creating better special effects. We should add that while it is fairly straightforward to describe increasing the number of rockets in a fireworks display as making the fireworks display “larger,” the same cannot quite be said for copyright goods. Better talent or higher-quality special effects does not mean that there is “more” movie or that the movie is “larger.” Really, it changes the artistic content of the movie and, to this extent, makes it a different movie. This debate aside, once the “size” of a copyright good—or indeed any public good—is determined, everyone must necessarily consume a public good of exactly that size.¹⁹

The nonrivalry of the benefits of public goods fundamentally changes the optimality conditions. Welfare maximization requires that production be increased until the benefits from doing so no longer exceed the costs. When goods are nonrival, the incremental benefit of expanding the size of the public good is represented by the sum of the marginal benefits of all consumers (i.e., $\sum_i MRS_i = MRT$).²⁰

Although every consumer necessarily consumes a public good of the same size, different consumers may derive different levels of utility from doing so. For example, some people enjoy fireworks displays more, while some enjoy them less. When everyone necessarily consumes the same quantity, rather than pay the same price and signal the intensity of their preferences by purchasing different quantities, as is the case with private goods, purchasers of public goods consume the

¹⁹ Some might assert that the fact that agents purchase different numbers of copies of particular works undercuts the claim that all agents consume a public good of the same size. This argument ignores the fact that copyright only protects the creative expression embodied in a work and “is distinct from ownership of any material object in which the work is embodied.” 17 U.S.C. § 202 (2006). Thus, no matter how many physical copies of a DVD that an agent purchases, the size of the creative work the agent consumes (measured in terms of the amount of resources used to produce the work) is exactly the same. See Yoo, *supra* note 4, at 667 (showing how copyright distinguishes between the creative aspects of works, which are both indivisible and protected by copyright, and the material objects in which the work is embodied, which are divisible and not protected by copyright).

²⁰ This equation is known as the Samuelson condition. It follows from the insight that the demand curve for private goods is derived by summing each individual’s demand curve horizontally, while the demand curve for public goods is derived by summing each individual’s demand curve vertically. See Yoo, *supra* note 4, at 662-71 (outlining the role of the Samuelson condition in public goods theory).

same quantity and can only signal the intensity of their preferences by paying *different prices*.

The key implication of the Samuelson condition is that price discrimination is a necessary condition for optimal production of public goods. Consider the following example: If a fireworks firm asked two individuals how much they would pay to see five rockets, the individual who really liked fireworks might be willing to pay a price of \$3 per rocket, whereas the other individual might only be willing to pay \$2 per rocket. Since rockets are nonrival, the firm need only fire five rockets altogether to satisfy both individuals, and the aggregate benefit of firing off the five rockets is \$5 per rocket. Under uniform pricing, a profit maximizer would set the price at \$2 per rocket and earn a revenue of \$4 per rocket. If the cost per rocket were \$4 or less, the fireworks display would enhance welfare and thus would occur. If, however, the cost per rocket were between \$4 and \$5, uniform pricing would not provide sufficient revenue to cover the cost of the fireworks display even though its benefits would exceed its costs. On the other hand, if price discrimination were permitted, the fireworks operator could charge the first consumer \$3 per rocket and the second consumer \$2 per rocket, which in turn would generate sufficient revenue for the welfare-generating fireworks display to occur. Thus, preventing price discrimination causes a systematic underproduction of public goods, while facilitating price discrimination can help promote the creation of more welfare-enhancing public goods.

A more fundamental problem posed by public goods is that there is no practical, real-world, incentive-compatible way to induce consumers to use prices to signal the intensity of their preferences. The nonrivalry of the benefit combined with producers' inability to verify the true marginal benefit each consumer derives gives consumers the strategic incentive to understate the intensity of their preferences in the hopes that some other consumer will incur the costs needed to increase the size of the public good. In fact, reporting no benefit and free riding on the public contributions of others appears to be a dominant strategy in most cases.²¹

²¹ See Samuelson, *supra* note 5, at 334, 336 (noting that because public goods "*simultaneously* enter into many persons' indifference curves," rational actors will "hide their desires for public goods" and "dissemble, trying to mask [their] preference[s] for the public goods and to engage in other game-strategy maneuvers which, when all do them, will necessarily involve deadweight loss to society"); Samuelson, *supra* note 3, at 388-89 ("[I]t is in the selfish interest of each person to give *false* signals, to pretend to have less interest in a given collective consumption activity than he really has . . .").

Stated somewhat more formally, with uniform pricing it is economically rational for any particular individual to increase her consumption only until the marginal benefit she derives (i.e., her individual marginal rate of substitution) equals the marginal cost (i.e., the marginal rate of transformation). Optimality, however, requires that the level of production of a public good be based not on any one individual's marginal rate of substitution, but rather on the sum of the marginal rates of substitution of every individual. The lack of incentive for any individual to take into account the marginal benefits derived by others represents the true source of systematic bias toward underproduction associated with public goods. Indeed, as Samuelson noted, this bias will exist even if the public good is fully excludable and is priced at marginal cost.²²

II. A FORMAL MODEL OF PURE PUBLIC GOODS

The insights of the foregoing discussion can be captured more formally by the following model: Assume the economy consists of one private good x ; N public goods y_1, \dots, y_N ; I individual agents with utility functions $U^i(x^i, y_1^i, \dots, y_N^i)$ and endowment of private good $\Omega^i \in \mathfrak{R}^1$ for $i = 1, \dots, I$; and one firm with production function $F(x, y_1, \dots, y_n) = 0$. The social planner's problem is to choose an allocation of public goods and private goods for each agent that maximizes a welfare function that weighs the implied utility levels received

Interestingly, even if each consumer takes the amounts that others will contribute for a public good as given, it will sometimes be rational for some agents to contribute some positive amount toward the provision of the public good. These contributions, while positive, will be well below socially optimal levels. See Theodore Bergstrom, Lawrence Blume & Hal Varian, *On the Private Provision of Public Goods*, 29 J. PUB. ECON. 25, 25 (1986) (noting that while instances of voluntary contributions occur, pure public goods would be undersupplied by these alone); Peter G. Warr, *The Private Provision of a Public Good Is Independent of the Distribution of Income*, 13 ECON. LETTERS 207, 207 (1983) ("When a public good is voluntarily provided, the level of its provision will typically be sub-optimal from a welfare standpoint."). As a result, the literature sometimes refers to this effect as "easy riding" rather than "free riding." See, e.g., CORNES & SANDLER, *supra* note 9, at 30 (noting that the fact that agents do make some contribution makes "easy riding" a better description of the suboptimality associated with public goods than "free riding").

²² See Samuelson, *supra* note 5, at 335 (using the example of subscription television broadcasting to demonstrate that converting a public good into a private good does not result in optimal production, as returns in consumption are not a scaled function of cost; rather, the marginal cost of one additional individual tuning in is zero).

by different agents and combines them into a social aggregate. The social planner's problem is thus:

$$\begin{aligned} & \text{Max}_{x^1, \dots, x^I, y_1^1, \dots, y_N^1, y_1^2, \dots, y_N^2, \dots, y_1^I, \dots, y_N^I} W(U^1, \dots, U^I) \text{ such that} \\ & x = \sum (x^i - \Omega^i) \\ & y_n = y_n^i \forall i = 1, \dots, I, \text{ and } n = 1, \dots, N \\ & F(x, y_1, \dots, y_N) = 0 \end{aligned}$$

The first condition requires that the net production equal the net consumption of private good. The second condition says that all agents consume the total amount produced of each public good. The third condition establishes that the production plan is feasible.

The production function gives the net private good surplus or deficit associated with any given production plan. Typically these plans involve negative private good components (since the private good is an input) and positive levels of public goods (since these are outputs). If the production function takes a value of zero, then the plan is feasible and the requested inputs are sufficient to produce the planned outputs.

The private good is purely rival. Thus, the sum of net consumption over all the agents (what each consumes less what each is endowed with) must equal exactly the net production of each private good. Since net production will generally be negative, the average agent will consume less private good than that with which the agent was originally endowed. The public goods, on the other hand, are purely nonrival. Thus, each agent consumes exactly the total production of each public good.

This approach requires the statement of a welfare function that weighs the well-being of each agent against that of all the others. Fortunately, it turns out that, for the major result, this interpersonal utility comparison is not necessary. We set this up as a constrained optimization using the method of the Lagrange multiplier. After some simple manipulation of the resulting first-order conditions we get the following:

$$\sum \frac{\frac{\partial U^i}{\partial y_n^i}}{\frac{\partial U^i}{\partial x^i}} = \frac{\frac{\partial F}{\partial y_n}}{\frac{\partial F}{\partial x}} \rightarrow \sum MRS_{y_n, x}^i = MRT_{y_n, x} \quad \forall n = 1, \dots, N$$

$$\frac{\partial W}{\partial U_i} \frac{\partial U^i}{\partial x^i} = \frac{\partial W}{\partial U_j} \frac{\partial U^j}{\partial x^j} \quad \forall i, j = 1, \dots, I$$

The first condition is the Samuelson condition. It requires that, for a nonrival good, production should increase until the sum of the marginal benefits (or the marginal rates of substitution of public for private good) across all agents equals the marginal cost (or the marginal rate of transformation of public for private good). The Samuelson rule is a necessary condition that must be satisfied by any Pareto-optimal allocation.

To ensure full social optimality, we must also consider the second condition. This is quite straightforward:

$$\frac{\partial W}{\partial U_i}$$

is the incremental social welfare that giving one more unit of utility to agent i produces for the society. On the other hand,

$$\frac{\partial U^i}{\partial x^i}$$

is the incremental utility that giving one more unit of private good to agent i produces. Thus,

$$\frac{\partial W}{\partial U_i} \frac{\partial U^i}{\partial x^i}$$

is the incremental social welfare produced by giving agent i one more unit of private good.

The equation, therefore, says to distribute the private good across agents such that the social welfare produced on the margin from pri-

vate good consumption is equated over all agents. Of course, there is no need for a similar distributional equation for public goods since all agents consume the total level of public good produced. In other words, structurally, there is never a distributional issue for public goods.

Achieving a full social optimum requires stating a social welfare function that explicitly weighs one agent's utility against the others'. This is a difficult exercise to imagine in the real world, but fortunately it is not important for our current objectives. The Samuelson condition holds independently of the welfare function or particular cardinalization of preferences chosen when the utility functions are stated. Of course, it also holds in more complex economies with many firms and many private goods.

A. *Perfect Price Discrimination*

We now turn to the more interesting question of how to implement a Pareto-optimal allocation using a market-based approach. Samuelson proposed an adaptation (based on the work of Erik Lindahl) of the standard notion of competitive equilibrium to an economy that includes public good.²³ The notion is quite straightforward. As discussed earlier, for private goods, there is a single price that is common to all agents and all firms. For public goods, however, each agent must have a personalized price that reflects her own personal marginal benefit from the public good.

These "Lindahl prices" decentralize the production and consumption of public goods as follows: Suppose that the optimal quantity of a public good was y (that is, y is the quantity of public good that satisfies the Samuelson conditions above). Then each agent would be given a personalized price q^i for the public good and would take this price as given. Just as agents do for private goods, each consumer would then demand a quantity of the public good such that her own marginal benefit equals q^i (which is her personalized marginal cost). Firms, on the other hand, are offered a production price of $\sum q^i = q$ and, taking this price as given, produce until the marginal cost of production equals the price for which they can sell the public good. One can immediately see the connection between Lindahl prices and the Samuelson condition.

²³ Samuelson, *supra* note 3, at 387-88.

To be a bit more formal, a Lindahl equilibrium consists of a set of prices $(p, q_1^1, \dots, q_N^1, q_1^2, \dots, q_N^2, \dots, q_1^I, \dots, q_N^I) \in \mathfrak{R}^1 \times \mathfrak{R}^{IN}$ and an allocation $(x^1, \dots, x^I, y_1, \dots, y_N) \in \mathfrak{R}^I \times \mathfrak{R}^N$ such that

1. $\forall i = 1, \dots, I$ $p \Omega^i \geq p x^i + q_n^i y_n^i$, and $\forall (\bar{x}^i, \bar{y}^i)$,
such that $p \Omega^i \geq p \bar{x}^i + q_n^i \bar{y}_n^i$, $U^i(x^i, y^i) > U^i(\bar{x}^i, \bar{y}^i)$
2. $F(x, y) = 0$, and $\forall (\bar{x}, \bar{y})$ such that $F(\bar{x}, \bar{y}) = 0$.
 $p x + \sum \sum q_n^i y_n \geq p \bar{x} + \sum \sum q_n^i \bar{y}_n$
3. $x = \sum_i (x^i - \Omega^i)$ and $\forall i = 1, \dots, I$ and $n = 1, \dots, N$ $y_n = y_n^i$

The first condition says that when each agent takes her personalized public good price as given, she chooses an affordable level of public and private good that maximizes her utility within her budget constraint. The second condition says that when the firm takes the sum of agents' personalized public goods prices as the net price for which it can sell its production of public goods, the firm chooses a feasible production plan that is profit maximizing. The final condition says that the supply and demand for private good are equal and that the amount of any given public good demanded by the agents is the same as the total amount produced.

The key conclusion from this analysis is that price discrimination is necessary for the efficient provision of public goods and, by extension, copyright goods. More specifically, if we declare by law that a single price prevail in the market, we will see losses of two separate types. First, we have the familiar loss due to low-value users finding the price of the copyright good too high. These users will choose not to buy the copyright good even though social welfare would be strictly improved by their doing so. Since the good is nonrival, the consumption benefits they would receive are a pure bonus to social welfare and impose no additional costs on copyright providers or other agents. Second, if some users who would be willing to make a positive contribution to help support copyright good-production are discouraged from doing so because the uniform price is too high, many marginal copyright goods will not be able to cover production costs and so will

not be produced at all. This results in a loss of both consumer and producer surplus.

In other words, even though copyright gives providers a degree of exclusivity, price discrimination is not per se bad for society. Price discrimination both expands access and increases production of copyright goods. Thus, the presumption may even go the other way. Moreover, the Lindahl/Samuelson approach implies that optimal provision depends on the producers' ability to extract the marginal benefit enjoyed by all consumers. Thus, price discrimination should be facilitated with respect to all consumers, not just low-value consumers who would be excluded by uniform pricing.²⁴

The Lindahl/Samuelson conception of efficient public good provision has two other features worth mentioning. First, even though Lindahl prices are personalized (and thus Lindahl pricing appears to be a kind of first-degree price discrimination), consumers still retain surplus. This is because Lindahl prices equal the *marginal* surplus gained by expanding the amount of the public good, rather than the total surplus generated by the public good. In other words, the first units of public goods purchased at the personalized Lindahl price by a consumer yield more consumption benefit than they cost. It is only the last unit of a public good that is priced at marginal benefit. As a result, Lindahlian pricing is less hostile to consumer welfare than traditional first-degree price discrimination.

Second, Lindahlian price discrimination falls on the same difficulty as other forms of price discrimination: identifying each consumer's willingness to pay. Thus, in proposing this solution, we are making the point that in the best of all possible worlds, where identification is feasible, price discrimination would result in the highest possible social welfare. In the real world, we must consider second-best solutions such as those we discuss below.

As we mentioned earlier, one objection that might be made to the analysis above is that it is difficult to interpret authors as being able to vary the size of the copyright goods they create in response to consumer demand. A book is a book, and a song is a song. In other words, they are discrete products provided and consumed in an all-or-

²⁴ Cf. Yoo, *supra* note 4, at 674 (explaining that to satisfy the Samuelson condition the producer "must be able to price discriminate over the entire range of output" and calling this a "dramatic expansion of the range over which price discrimination is important").

nothing way. On the other hand, they are still public goods since consumption is nonrival.

What are the Samuelson conditions in this case? Suppose that we have an arbitrary set of potential copyright goods that could be produced. Each of these goods produces a personalized benefit for any given agent and has a cost of production associated with it. Assume for simplicity that these goods are neither complements nor substitutes for one another. The economy consists of a set of J potential copyright goods, where $j \in \{1, \dots, J\}$ denotes a specific good; and I individual agents, where $i \in \{1, \dots, I\}$. The benefit to agent i of consuming good j is given by B_j^i and the cost of producing good j is given by C_j .

The highest possible overall welfare is achieved in this model if each agent is charged B_j^i for the use of any copyright good and if, based on these prices, copyright good producers make profit-maximizing provision decisions. It is interesting that this is a sufficient, but not a necessary, condition for Pareto optimality, which is the inverse of what we found for the Samuelson conditions in the case of divisible public goods given earlier.

Observe that copyright producers would agree to provide good $j \in \{1, \dots, J\}$ if and only if $\sum B_j^i \geq C_j$. This means that all goods for which social benefits exceed social costs will be produced. On the other hand, all agents will choose to consume all produced goods at these prices since their private benefits are at least as big as the cost of consuming the goods. It follows both that we have the optimal number of goods produced and that no agent who would get any benefit from consuming a copyright good is denied access.

A problem here, however, is that there is no consumer surplus. While this does not reduce efficiency, it may be undesirable from a distributional standpoint. To address this, policymakers could place an upper limit on the price charged for access to a copyright good. This might be something like a compulsory license. While this would leave the upper-end consumers some surplus, it would also reduce the number of copyright goods offered. Thus, we see the traditional balance between equity and efficiency that is familiar in economics. However, regardless of the restrictions on the highest price allowed, producers would still prefer to offer low-value consumers lower usage prices. This only increases firms' revenue. Allowing this kind of low-end price discrimination increases social welfare both by increas-

ing the number of users who get to enjoy a given copyright good and by making copyright goods more profitable and so more numerous. Thus, it is hard to see how a good policy case could be made to foreclose this kind of low-end price discrimination.

It bears reemphasizing that this type of price discrimination represents only one of several ways for markets to produce indivisible public goods efficiently. As we noted earlier, allowing the producer to capture all the available surplus is a sufficient condition for Pareto optimality, but it is not a necessary one. This implies that other pricing solutions may exist that permit consumers to retain some surplus while also supporting the efficient market provision of the public good.²⁵

B. *Imperfect Price Discrimination*

The foregoing example is limited by the fact that firms generally will not be able to fully identify each consumer's willingness to pay. We are therefore in a second-best world. What should be the policy considerations in this case? More specifically, how intensive should copyright be when the objective is to maximize social welfare? We might allow a copyright holder to choose a single price (and less intensive copyright) or allow second-degree price discrimination with two prices for different classes of consumers (and more intensive copyright). One might imagine that making copyrights more intensive would simply permit the copyright holder to increase her exploitation of consumers. However, the structure of markets for copyright goods has a feature that makes it particularly likely that consumers will actually benefit. Specifically, by allowing two prices, the use of the copyright good can be extended to additional consumers who would have been priced out of the market under the less intensive single-price copyright regime.

The following example illustrates this point. To keep matters simple, we consider a demand system with five different classes of con-

²⁵ For one perspective on this problem, see John P. Conley & Paul J. Healy, *Public Goods, Bounded Attention Spans and Equilibrium in the Internet Economy* (Mar. 2008) (unpublished manuscript, on file with author), available at <http://www.vanderbilt.edu/Econ/jpconley/documents/PE25-internet-attention-infinite-public-goods.pdf>. The authors model a public goods economy with a continuum of agents in which limitations in attention span cause them to consume finite amounts of a finite number of public goods. They show that while price systems that decentralize efficient public good provision and allow agents to retain surplus do exist, the complexity of the required price system is extreme.

sumers. This is not essential for the example to work, but it simplifies the calculations. Assume the demand system is as follows:

Table 1: Demand System with Five Different Classes of Consumers

Number of Agents	Reservation Price
100	60
100	50
100	40
100	25
100	20

Suppose further that the copyright holder must choose one and only one price for her good. Thus, price discrimination is prohibited. The uniform pricing solution is as follows:

Table 2: Uniform Pricing Solution

Price	Profit	Consumer Surplus
60	6,000	0
50	10,000	1,000
40	12,000	3,000
25	10,000	7,500
20	10,000	9,500

The profit-maximizing decision is for a firm to set a price of 40 and obtain a profit of 12,000. This in turn yields a consumer surplus of 3,000.

Now assume that the copyright holder is able to perfectly identify the reservation price of consumers and that resale is impossible. Also suppose that the copyright holder is allowed to employ third-degree price discrimination and set different prices in different markets. Table 3 shows every possible pair of prices and the associated profit and consumer surplus.

Table 3: Every Possible Pair of Prices and Their Associated Profit and Consumer Surplus

Price	Profit	Consumer Surplus
60/50	11,000	0
60/40	14,000	1,000
60/25	13,500	4,000
60/20	12,000	5,500
50/40	14,000	1,000
50/25	15,000	2,500
50/20	16,000	6,000
40/25	14,500	3,000
40/20	16,000	5,500
25/20	12,000	7,500

There are two profit-maximizing choices: 50/20 and 40/20. Both yield 16,000 in profits and give 6,000 and 5,500 in consumer surplus, respectively. Both of these give higher profits and higher consumer surplus than uniform pricing. Thus, allowing a more intensive copyright by facilitating price discrimination has actually improved total welfare for all classes of agents in the economy.

What this example is meant to show is that, in a more realistic setting, it is far from obvious that policymakers should discourage price discrimination. For copyright goods (which are really discretely provided public goods), price discrimination may very well improve welfare. It should certainly not be judged per se bad.

III. SPATIAL COMPETITION/IMPURE PUBLIC GOODS

One can think of copyright goods as being bundles of hedonic characteristics.²⁶ For example, many science-fiction films are similar in many ways (flying spacecraft, laser weapons, etc.) but different in

²⁶ For the seminal works, see Kelvin J. Lancaster, *A New Approach to Consumer Theory*, 74 J. POL. ECON. 132, 134 (1966), which offers a model of consumer behavior that treats goods as collections of characteristics from which each consumer derives a different level of utility, and Sherwin Rosen, *Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition*, 82 J. POL. ECON. 34, 34 (1974), which sets forth a model of product differentiation based on the hedonic assumption that "goods are valued for their utility-bearing attributes or characteristics." The discussion that follows expands on earlier work. See Yoo, *supra* note 4, at 687-706 (applying the insights of product differentiation to copyright).

others (characters, settings, plot, etc.). J.R.R. Tolkien's classic trilogy, *The Lord of the Rings*, inspired a whole host of fantasy novels that draw on similar imagery (elves, dragons, wizards, magic, etc.), but each has its own distinctive characteristics. The *Charlie Bone* series of books describes a castle-based school where children learn how to do magic, but it does so with its own unique characters and a much darker tone than the *Harry Potter* series.

Exactly how to rank, order, and compare such hedonic characteristics is an interesting and difficult question in both economics and law. Clearly goods like these are imperfect substitutes and may sometimes be so similar as to produce violations of copyright.²⁷ Producers of such copyright goods often argue that copyright protection should be as extensive as possible. They also engage in a wide array of price discriminatory practices (such as different release dates in different countries and at different prices, delayed release of DVDs and sales to cable outlets, etc.). The policy question is, how do these practices affect consumer surplus, total welfare, the variety of copyright offerings, and access to those works? The answers are not as clear-cut as one might think.

The potential beneficial impact of price discrimination is easily illustrated within the framework of spatial competition pioneered by Harold Hotelling.²⁸ Hotelling's approach assumes that goods are located at points on the unit interval (i.e., a line one unit long) and that agents are also located in a uniformly distributed manner along a unit line with a uniform density of one per unit length. In the context of local public goods, agents have a physical location and prefer to visit the public facility closest to them to minimize travel costs. Generalizing this idea to a space of hedonic characteristics means interpreting locations as embodying different consumption attributes of goods. In this context, an agent's location is equivalent to her most preferred good characteristic.

Mapping public goods into a hedonic space allows us to begin to understand how consumers choose among different offerings and how public good producers compete with one another for consumers. Doing so does not turn the relevant goods into private goods. Nothing about allowing for competition among different product charac-

²⁷ See, e.g., *Twentieth Century-Fox Film Corp. v. MCA, Inc.*, 715 F.2d 1327, 1329 (9th Cir. 1983) (holding that plot similarities between *Battlestar Galactica* and *Star Wars* raised genuine issues of material fact sufficient to preclude summary judgment).

²⁸ See Harold Hotelling, *Stability in Competition*, 39 *ECON. J.* 41 (1929) (setting forth the seminal model of spatial competition).

teristics alters the fact that each agent consumes the entire industry output. Thus, optimality still depends on satisfying the Samuelson condition.

The existence of hedonic characteristics introduces an additional optimality criterion that strikes a balance between the increase in fixed costs needed to establish another producer and the increase in utility from enabling agents to consume goods that fit better with their preferences. The standard results of hedonic models are that, in equilibrium, markets may produce too many or too few public goods. Thus, rather than supporting the simple policy inference that markets tend to produce too few public goods, these models suggest the absence of any systematic bias pointing in either direction. These models also suggest that price discrimination may bring the equilibrium closer to optimality, which raises serious questions about the reflexive hostility toward practices that tend to facilitate price discrimination exhibited by much of the current literature.²⁹

The underlying intuitions are captured by the following example, in which three imperfectly substitutable copyright goods, *A*, *B*, and *C*, can be produced. There are three types of consumers (α , β , and γ) who have single-peaked preferences over these goods. There are five individuals of each type and the surplus/willingness to pay of each type is given in the table below.

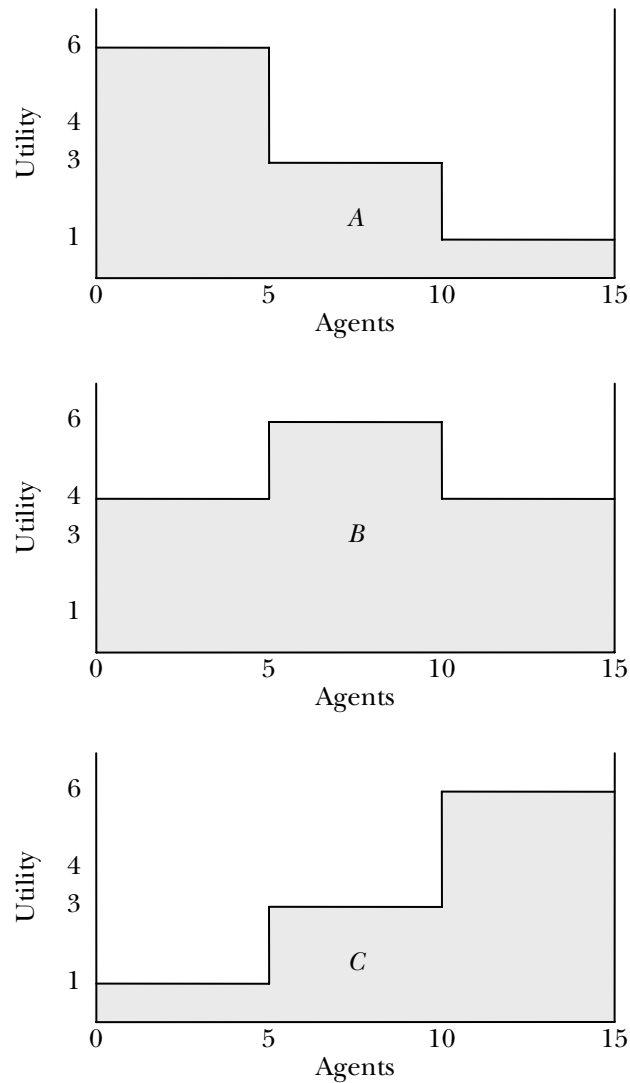
Table 4: Single-Peaked Preferences of Three Different Types of Consumers

	Willingness to Pay (WTP)		
	Good A	Good B	Good C
Type α (Agents 1-5)	6	4	1
Type β (Agents 6-10)	3	6	3
Type γ (Agents 11-15)	1	4	6

²⁹ See, e.g., Joseph E. Stiglitz, *The Theory of Local Public Goods*, in *THE ECONOMICS OF PUBLIC SERVICES* 274, 309-12 (Martin S. Feldstein & Robert P. Inman eds., 1977) (showing how a Hotelling linear spatial model can lead to the socially optimal level of production as well as either excess or insufficient entry).

The figure below illustrates these preferences in the more familiar Hotelling/Stiglitz framework. When an agent's location is equivalent to her most-preferred-good characteristic, she loses utility if she "travels" away from her most-preferred-taste location and consumes another good instead.

Figure 1: Illustration of Consumer Preferences in the Hotelling/Stiglitz Framework



Of course, each good has a production cost. This is given in the table below.

Table 5: Production Cost of Each Good

Cost		
Good A	Good B	Good C
25	40	25

Suppose that price discrimination were illegal or infeasible. This would mean that the provider of each good must choose one price and offer this to all agents. Also suppose that copyright scope were sufficiently limited such that none of these goods infringed on one another.

To understand how the market would work in this case, suppose at first that only good *B* were provided. If provider *B* charged 6, only type β 's would buy the good, revenue would be 30, and profit would be -10. On the other hand, charging 3 would mean that all agents would buy the good, revenue would be 45, and profit would be 5. This is therefore the profit-maximizing outcome.

Observe, however, that provider *A* could enter and charge a price of 5.5. At this price, type α 's would choose to purchase good *A* and give provider *A* a total revenue of 27.5 and a profit of 2.5. To lure type α 's back to good *B*, provider *B* would have to drop her price to 2.5 in order to match the consumer surplus of 0.5 that the α -type agents currently received from consuming good *A* at a price of 5.5. But this would only generate revenue of $15 \times 2.5 = 37.5$ for producer *B* and so would not cover her costs of 40. Of course producer *C* could also enter, offer a price of 5.5, and attract all type γ 's.

What we see is that there is nothing that provider *B* can do regarding price to deter entry of goods *A* and *C*. However, surplus with good *B* is $5(3 - 3) + 5(6 - 3) + 5(3 - 3) + 15(3) - 40 = 20$.³⁰ With goods *A* and *B* produced, surplus is $5(6 - 5.5) + 5(5.5) - 25 + 5(6 - 5.5) + 5(5.5) - 25 = 10$.

³⁰ Note that $5(3 - 3) + 5(6 - 3) + 5(3 - 3) + 15(3) - 40 = 5CS_{\alpha} + 5CS_{\beta} + 5CS_{\gamma} + TR - TC$, where CS_x is the consumer surplus of x , and TR and TC are total revenue and total cost, respectively.

What does this illustrate? Surplus is highest when only good *B* is provided. Surplus goes down when goods *A* and *C* are also provided. More agents consume a copyright good when only good *B* is provided. All agents buy good *B*, but only α 's and γ 's consume copyright goods when *A* and *C* are provided.

Thus, a market with free entry can actually generate excess entry both from a standpoint of total surplus produced and total consumers served. The existence of equilibria with excess entry thus undercuts any policies based on the broad inference that markets systematically produce too few public goods. Moreover, in markets that reach equilibrium with too many goods, measures that increase access to works may actually promote optimal incentives, since the resulting reduction in incentives may actually promote efficiency by compensating for the market's tendency toward excess entry. When that is the case, the access/incentives dichotomy that frames most economic analyses of copyright does not exist.

This is not to say that decreasing copyright's scope will always be welfare enhancing. Suppose that *A* and *C* were copyright works that were derivative of good *B*. Then we see the counterintuitive outcome that making copyright more extensive can increase access to copyrighted works, lower prices, and raise overall welfare.³¹

This example is predicated on one price prevailing in the market place. Let us suppose instead that we permitted perfect price discrimination. First, consider the planner's problem. The table below calculates the net social surplus for each possible combination of goods.

³¹ This example underscores how the hedonic approach provides a basis for distinguishing among different ways copyright can be strengthened or weakened. Copyright can be strengthened either by adding more surplus-generating activities within its scope (making the right larger), by facilitating price discrimination (making the right more intense), or by reducing the degree of similarity required in order to support a finding of infringement (making the right broader). Thus, when faced with equilibria in which there is insufficient entry, the optimal approach might be to strengthen copyright along the first two dimensions (making it larger and more intense) while weakening it along the third (making it narrower). See Yoo, *supra* note 10, at 264-76 (finding that, under the differentiated-products approach, copyright protection that is both large and intense is most conducive to a competitive market for copyright).

Table 6: Net Social Surplus for Each Possible Combination of Goods

Goods	Total Benefit	Total Cost	Net Social Surplus
<i>A</i>	55	25	30
<i>B</i>	60	40	20
<i>C</i>	55	25	30
<i>A + B</i>	80	65	15
<i>B + C</i>	80	65	15
<i>A + C</i>	80	50	30
<i>A + B + C</i>	90	90	0

Suppose that we gave copyright to the whole range of products to the producer of good *A* (or *C*) and permitted price discrimination. Even if good *A* charged the maximum price (WTP), the surplus would be 30 and all agents would consume the good.

Suppose instead that we did not grant extensive copyright but allowed price discrimination. Suppose we started with producer *A* charging each agent her WTP. Provider *C* could enter and undercut these prices. The lowest price producer *C* could set for type α 's would be 0. This would give α 's one unit of utility. Thus, producer *A* would have to set a price of 5 to retain type α 's. What about type β 's? Although we have identification and we assume that no resale is possible, we do not have monopoly. Thus, *A* and *C* would have to compete for type β 's by undercutting one another's price. Since type β 's would be indifferent between *A* and *C*, the only equilibrium price would be 0. How type β 's agents split between *A* and *C* is not important. Finally (and symmetrically), provider *C* could sustain a price of 5 for type γ 's. In summary, α 's pay 5 and consume good *A*, γ 's pay 5 and consume good *C*, and β 's pay 0 and consume either good. Thus, by allowing price discrimination but only narrow copyright, we would obtain the highest surplus, the lowest consumer prices, and largest market served.

Could provider *B* enter at these prices? The highest price it could charge α 's and γ 's would be 2 (and still give them a surplus of 1). The highest price provider *B* could charge β 's would be 3 (and still leave

2009]

Nonrivalry and Price Discrimination

127

them with a surplus of 3). Thus, revenue would be 35, which is not enough to pay for the project, and so deters this entry.

The following table summarizes the conclusions we get from this example:

Table 7: The Result of Different Levels of Copyright Protection and the Presence of Price Discrimination

Copyright Protection	Narrow ©		Extensive ©	
	No PD	No PD	With PD	With PD
Number of Copyright Goods	2	1	2	1
Consumers Served	2/3	All	All	All
Consumer Surplus	5	15	30	0
Total Net Welfare	10	20	30	30

What we can observe generally is that price discrimination plus extensive monopoly will always increase access to copyright goods, although it may lower consumer surplus. On the other hand, it may increase or decrease overall welfare.

Allowing price discrimination but then decreasing the extent of monopoly tends to lower prices and leave more surplus for consumers. All else being equal, price discrimination always increases access to copyright goods. However, without a reasonably extensive monopoly, price competition may make it unprofitable to provide socially beneficial goods. Thus, we can get underprovision of copyright goods if the monopoly is not sufficiently extensive. (To see this, make the trivial change of raising the cost of goods *A* and *C* from 25 to 26 in the example above. The market will only sustain one incumbent, and the problem becomes a standard contestable-market issue.)

Disallowing price discrimination can make possible the excessive provision of copyright goods, but, surprisingly, can also result in more limited access and diminished welfare. This is because forcing providers to charge a single price prevents providers of broadly popular goods from offloading a larger share of the production costs on high-demand consumers. This allows producers of more narrowly focused goods to draw off segments of the consumer base to the point that the broadly popular good is not profitable. We end up with a few high-value niche markets being served but the broad consumer base being ignored.

Acknowledging the possibility of demand interactions among public goods thus belies the simplicity of the policy inferences usually drawn by copyright scholars applying public goods theory. The systematic bias toward producing too few public goods disappears. In addition, price discrimination may well be welfare enhancing. That said, results of this type depend on a number of assumptions. For example, the simple Hotelling models discussed in this Article are “discrete-choice” models that presume that agents only consume the good located closest to them. In addition, they assume that agents’ preferences are single peaked and can be ordered into a linear spectrum. Furthermore, they presume that the hedonic space consists of a single dimension. These models can be extended in ways that permit consumption of multiple goods, multi-peaked preferences, and multidimensionality, but the added complexity prevents the derivation of general results.

To say that these more general models do not yield general results is not to say that they yield no results at all, only that more parameterization and structure is needed on the nature of agents’ preferences and the structure of demand before any conclusions can be drawn. The fact intensiveness of the analysis that results when these restrictive assumptions associated with simple Hotelling models are relaxed undercuts further the propriety of basing copyright policy on simple policy inferences drawn from public goods theory.

CONCLUSION

The overwhelming majority of scholarly commentaries on the economics of copyright base their claims on the basic insights of public goods theory: that markets systematically produce too few public goods and underutilize those that are produced. Moreover, modeling nonrivalry as zero marginal cost leads the literature to frame copyright

as an inherent tradeoff between optimal incentives to create works and optimal access to works that are created, with any feasible solution presumed to be necessarily second-best in both aspects. The desire to promote access without increasing incentives has also led most commentators to support facilitating price discrimination only with respect to consumers of low-value uses, who would be inefficiently excluded from the market if price discrimination were prohibited altogether.

Framing the issues in this manner has caused the existing scholarship to overlook the economic features that give public goods their distinctive quality. In particular, the copyright literature has almost entirely failed to appreciate the importance of the Samuelson condition, which remains one of the most fundamental insights of public goods theory. In particular, the Samuelson condition implies that price discrimination is not only not necessarily harmful but also that, absent direct government subsidies, it is a necessary condition for efficient provision of public goods. In addition, it has implications for the scope of price discrimination, showing that it is important with respect to all consumers of public goods and not just consumers of low-value uses.

To the extent that public goods are divisible, producers need only capture the sum of the marginal benefits that consumers derive from a public good. This ensures that consumers retain some portion of the surplus while also providing a metric for determining the optimal level of price discrimination. If public goods are indivisible, as is typically the case with creative works, perfect price discrimination in which producers appropriate the entirety of the available surplus is a sufficient condition for optimality, although the fact that it is not a necessary condition implies the existence of alternative equilibria in which consumers do retain surplus. These alternatives presume that the producer is able to charge individualized prices calibrated to the precise benefit each consumer derives from the public good. In the more realistic setting of imperfect price discrimination, it is still quite plausible that facilitating price discrimination would cause both consumer surplus and total welfare to increase.

The results become even more interesting when Samuelson's initial approach, in which the set of public goods is determined exogenously and has no demand interactions, is broadened to allow for demand interactions and for the number of public goods to be determined endogenously. Under such models, the systematic bias toward underproduction of public goods disappears entirely. Further-

more, the models suggest that disallowing price discrimination and restricting the breadth of copyright protection may actually reduce both consumer surplus and total welfare and may induce excess entry. Although relaxing some of the strict assumptions on which the classical models are based may make for a more realistic approximation of markets for copyrighted works, this is achieved at the cost of adding complexity that makes general results even more difficult to obtain absent additional information on the precise structure of demand. The intricacy of the analysis only serves to further underscore the extent to which public goods theory fails to provide the type of simple policy inferences needed to provide the categorical conclusions that characterize so much of the copyright literature.