

Two (Lay) Dogmas on Externalities

Vaughn Bryan Baltzly

Texas State University

Department of Philosophy

baltzly@gmail.com

March 22, 2024

(forthcoming in *Public Choice*; this is the penultimate version)

ABSTRACT

I argue that much current thinking on externalities—at least among “lay political economists” (but even, on occasion, among professional economists)—is saddled with two analytical errors. The first is what I call *coextensivism*: the conflation of public goods and externalities. The second error is what I call *externality profligacy*: the conflation of economic and “social” externalities. The principal dangers presented by these two “dogmas on externalities” are that, while in their grips, we are under-disposed to seek negotiated, market-based solutions (of a broadly Coasean nature) to challenges posed by economic externalities, and over-disposed to seek coercive, state-based solutions (of a broadly Pigouvian nature) to challenges posed by social externalities.

With apologies to Willard Van Orman Quine,¹ we might say that much “lay discussion” of externalities has been conditioned in large part by two dogmas. Each involves a *conflation* of two things, where instead we should recognize a fundamental cleavage. The first dogma is what I call *coextensivism*: the conflation of public goods and externalities. More precisely, coextensivism is the view that, while there may be conceptual distinctions to be drawn here, the phenomena widely labelled “public goods” and “externalities” very often (or perhaps always) co-vary—they are, one might say, two sides of the same coin. The second dogma is what I call *externality profligacy*: the conflation of economic and social externalities. More precisely, the profligate identification of externalities involves subsuming externalities that are merely *social*—instances wherein we observe an “impact of one person’s actions on the well-being of a bystander” (as per Mankiw [2008: 200])—under the category of externalities that are *economic*, in the sense that they arise for parties that are “external to” other parties’ welfare-enhancing voluntary exchanges for mutual benefit. Both dogmas, I shall argue, are ill-founded. The dangers posed by their uncritical acceptance are, respectively, (1) the risk of occluding our view of various broadly Coasean solutions to various problems generated by economic externalities, and (2) the risk that we will bring economic and public policy tools (of broadly Pigouvian—but also sometimes even Coasean—natures) to bear on problems that are not properly *economic* or even *public*.

This paper proceeds as follows: The first section sets out, as a baseline, what I believe to be the proper and defensible understanding of the nature of externalities and their relation to public goods—a view I believe to be at least implicit in the work of most disciplined economic theorists. Since my focus

¹ From whose [1951] I have shamelessly appropriated, not only my title, but also my opening motif.

in this paper is two errors commonly committed by “the laity,” for purposes of contrast I shall call this the “Clerical View.” The second (and longest) section offers analyses of the dogma of coextensivism. In its first sub-section I attempt to steelman the coextensivist’s thesis—offering the strongest and most charitable rendering of the motivations and considerations that might have inclined so many to (at least tacit) acceptance of this view. In 2.2 I turn to critique, articulating three different ways in which coextensivism is mistaken. I then close this discussion, in section 2.3, with a brief sketch the sort of “Coasean clarity” that we might obtain when our thinking about externalities is excised of the coextensivist dogma. Section 3 addresses the dogma of externality profligacy: two sub-sections are devoted (respectively) to the dangers I there label as “over-Pigouving” and “over-Coaseing.” I then close this discussion, in section 3.3, with a brief sketch of a sort of “Coasean ideal of community”—an ideal that might become endangered, lest we clear away the dogma of externality profligacy. A brief fourth section concludes with some reflections about the nature of externalities and how we can classify them going forward.

Before beginning the detailed analysis of these dogmas, though, it’s useful to clarify the scope of my complaint. I have no doubt that many—perhaps even a large majority—of professional economists are perfectly clear on everything that follows, and entirely innocent of my allegations. Nevertheless, the complaints are worth articulating, for at least three reasons. First—and as the below discussion of (*inter alia*) Friedman, Cornes and Sandler, and Gruber reveals—there remains some (non-negligible) residue of influential working economists who apparently *aren’t* clear on these matters, and who *do* make the mistakes that are the concern of this paper.² Second, once we extend our consideration beyond the sphere of working professional economists, I suspect that there are significant swathes of the wider “political economy space” where said errors are rather more widespread. One encounters them more frequently, for instance, in the applied economics literature.³ And speaking personally for just one moment: my profession is academic philosophy; I have graduate-level training in both political philosophy and in public policy (and have taught graduate-level classes in both fields, as well as many undergraduate-level courses in the former), and I have seven years’ worth of professional experience in the U.S. federal government, split equally between the executive and the legislative branches. Though this is merely anecdotal, my strong impression is that in all these arenas—in academic philosophy, in the public policy academy, in the field of political science, and in the U.S. federal bureaucracy—the two dogmas here analyzed meet with widespread (albeit largely tacit) assent. The reader should therefore bear in mind that—even though I principally cite professional economists in what follows (for purposes related to the first of my three reasons, provided just above)—the lion’s share of my target lies outside academic economics. There is, if you will, something of a “folk political economy” in these professions that is prone to these errors, and it is this which is my primary concern. (In a similar context, political philosopher Michael Huemer [2013, xxvii-xxviii] has written of the need to engage, not only with the official written corpus of contemporary political theory, but also with its “oral tradition.” This impulse—to engage not only with formal, published work appearing in the relevant peer-reviewed journals, but

² In this context it may be worth noting that my present project has something of its own pedigree. Writing just over fifty years ago, S. E. Holtermann [1972, 78] complained of “a tendency in the literature on public goods to identify the two concepts [that is, public goods and externalities] with each other,” and advertised that “[t]he view taken here is that they are two distinct concepts”—even if it remains the case that “many externally produced commodities have the character of public goods.” I am clearly not alone, then, in taking note of this conflation, and evidently it was common enough among professional economists in the 1960’s and 1970’s. Half a century has elapsed since Holtermann published these remarks. While perhaps they had some effect in clearing up economists’ thinking on this score, apparently they did not suffice to completely banish the confusion—and hence the need, still, for the present paper.

³ See, e.g., the discussion of Costello and Kotchen [2022] and Jacobsen and LaRiviere [2017] in footnote 10 below.

also with more informal work appearing in alternative venues—is especially apt for our internet age, when so much of the scholarly community’s scholarly discussions transpire in “unofficial” forums, beyond just articles and books.) And this segues naturally into the third of my three reasons for engaging these dogmas. Even if a large majority of professional economists are perfectly clear of these dogmatic mistakes, the fact that the wider oral tradition and/or folk political economy betoken these confusions as widely as they do, is perhaps owing to the fact that economists have not sufficiently foregrounded the necessary clarity on these matters in their published work. In this regard: if this paper prompts even a handful of authors of economic textbooks (or the instructors who make use of extant texts) to take additional pains to make clear, to their readers and students, positions which (until now) they might have thought were too obvious to dwell upon—then this paper will have succeeded.

But first: if am I to depict the dogmas of coextensivism and externality profligacy as departures from some more proper (and professional) understanding of the relevant conceptual terrain, it behooves me to set forth a picture of what I think that terrain looks like. Though it’s perhaps not often spelled out quite this explicitly, what follows in Section 1 is a brief description of the framework that I presume informs the work of most careful professional economists working in the space of externalities and/or public goods.

1. The “Clerical View”

Economists study that subset of human interactions called “economic exchanges.” Not all human interactions are fruitfully modeled as *exchanges*, of course. (Sharing an afternoon stroll in the company of a loved one doesn’t seem to be, for instance.) And among those that are, not all of them are fruitfully modeled as *economic* exchanges. (When I return your favor—say, by watching your cats when you leave town next week, because you watered my plants while I was hospitalized last week—we’ve engaged in a sort of *quid pro quo*, but it doesn’t seem to merit the appellation of “economic exchange.”) The boundaries between non-exchanges and exchanges—and, with respect to the latter, between the non-economic and economic varieties—are of course fuzzy; we can’t draw them with any precision. Nevertheless, there seem to be clear cases on each side of the divide. (My previous two parentheticals sought to supply examples that were clearly on one side of each.) Furthermore, one always confronts the temptation to push the boundaries of “economic exchange” too far. When one succumbs to this temptation, one has overzealously expanded the notion of economic exchange, and has begun to model as *economic exchanges* the sorts of human interactions that, all things considered, are probably best exempted from analyses in terms of, e.g., *utility maximization*, *comparative advantage*, or (as I’ll describe more fully below) *Coasean bargaining*.

1.1 Externalities

One strand of theorizing about economic exchanges focuses on their “spillover” effects. On the picture presupposed by such theorizing, there are various entities that are party to a given economic exchange—call these the parties that are “internal to” that exchange—and then there are other (roughly speaking) “adjacent” parties—some of whom might be nevertheless indirectly impacted by said exchange. That is to say: there are parties who might be indirectly impacted by the exchange that transpires “directly” between (or among) the transaction’s internal parties. (Such indirect impacts include, of course, those generated by the production of the good or service itself, prior to its being exchanged.) Call these indirectly impacted parties those that are “external to” the exchange.

(It's tempting to call them "third parties," but this might misleadingly suggest that there are always only two parties internal to any given exchange.) And call the impacts to external parties the "external effects" of an exchange—or its "externalities" for short. These external effects might be beneficial to the external parties (in which case we have "positive externalities"), or they might be detrimental ("negative externalities").

In their modeling of externalities—whether formal or informal—economists always presuppose the existence of a *community* of some sort: the relevant population of (at least potentially-)impacted external parties. The size of this community might be relatively large or small, and its boundaries might be more or less determinate. Further, these specifications of size and boundary might be made relatively explicitly, or they might remain perfectly tacit. But it is always with respect to such specifications that the modelers (at least implicitly) characterize the *scope* of the externalities.⁴ That is, a given externality can be characterized (relative to some posited baseline community) as being of narrow scope—impacting only a small number of external parties, say (or, in the limit, only a lone individual)—or of moderate scope—impacting a non-negligible number of external parties—or of wide scope—impacting a majority or a supermajority of the community's members. And in some cases, we might even observe an external effect that is pervasive with respect to the members of the community. When the external effects of some economic transaction are thusly universal, we have what we might term a *ubiquitous* externality.

It's important that we be able to think clearly and systematically about externalities. For when we consider whether or not free exchange for mutual benefit increases social welfare, we have to attend, not only to such exchanges' internal effects (which are tautologically positive, else the exchange wouldn't transpire⁵), but also to their external effects (which may be sufficiently negative as to render the exchange welfare-diminishing on net). Efforts to, e.g., identify and measure various transactions' various external effects, and to seek clever and innovative ways to "internalize" these external costs and benefits—and/or to minimize the scope of negative externalities and maximize the scope of positive externalities—therefore represent fruitful ways of promoting social welfare. (This is because such efforts serve to render the mechanism of voluntary exchange for mutual benefit a progressively better and better engine for social welfare-maximization.) An important task of economic theorizing is to contribute to such efforts.

1.2 Public Goods

Another strand of theorizing about economic exchange focuses on a set of interesting properties that we can model certain goods and services as bearing—the properties of *non-excludability* and *non-rivalry*.⁶ Goods and services that can be fruitfully modeled as having the former property (at least to some considerable extent) are typically classified as common-pool resources, while those modeled as having a sufficient degree of the latter property are called club goods. Goods and services that bear *both* these properties are *public goods*, properly so called.⁷

⁴ For example, see the horizontal dimension of Fig. 1 in Paniagua and Rayamajhee [2023] elsewhere in this volume, which offers an explicit attempt to taxonomize externalities with respect to their scope in precisely this fashion.

⁵ Ignoring for now instances of non-voluntary exchange, as might be elicited by (e.g.) instances of coercion or deception or (depending on one's view on the nature of truly *voluntary* exchange) desperation.

⁶ Or, in a previous era, on "indivisibility" rather than non-rivalry.

⁷ The laity also exhibits a regrettable tendency to *improperly* deploy the term "public good." Oftentimes the phrase seems to be understood to mean roughly "anything that's generally in the public's interest"; other times it seems to be conflated

In their modeling of public goods—whether formal or informal—economists always presuppose the existence of a *community* of some sort—the relevant population of parties who cannot be excluded from consuming the good or service in question, once it is provided. The size of this community—that is to say, the scope of the relevant “public” for which a given non-excludable good constitutes a public good—might be relatively large or small, and its boundaries might be more or less determinate. These specifications of size and boundary might be made relatively explicitly and fairly precisely—talk of “global public goods” (such as reductions in atmospheric CO₂ concentrations) explicitly postulate a precisely planetary public (or, for all practical purposes, an *unbounded* public). Or they might tend towards the implicit and vague: discussions of, e.g., clean groundwater as a public good tacitly presuppose a fairly vague notion of a “region” or “local community” whose inhabitants utilize said water source. But it is always with respect to such specifications that the modelers (at least implicitly) characterize the *public* that is prerequisite to there being a “public” good.

It’s important that we be able to think clearly and systematically about public goods, for their central features (chiefly non-excludability) pose well-known challenges. Principally, the fact that non-excludable goods give rise to various market failures (via channels like the free-rider problem and the assurance problem: *collective action problems*, broadly speaking) poses a challenge to an overly-sanguine reliance on standard market mechanisms to maximize social welfare. For many (potential) welfare-enhancing public goods will go un- or under-provisioned, owing to these collective action problems—at least absent some clever, well-designed market manipulations or interventions (be they “statist” or otherwise).

* * * * *

The two strands of theorizing just described would appear to focus on wholly distinct categories of economic phenomena. The latter concerns itself with challenges that beset the provision of an interesting subset of goods and services—*viz.*, those that exhibit the properties of non-excludability and non-rivalry to any appreciable extent. The former concerns itself with the external impacts or effects of the production and/or exchange of goods and services that other parties (internal parties) find it welfare-enhancing to produce and exchange. One focuses on the (non-)production of goods and services themselves, while the other focuses on the (indirect) effect or impacts of the production and exchange of goods and services.⁸ Furthermore, the reasons we have to think carefully and systematically about attempts to “internalize externalities” (briefly summarized four paragraphs above) are not the same as the reasons we have to think carefully and systematically about public goods (briefly summarized in the foregoing paragraph). For one thing, the sorts of solutions (Coasean, Pigouvian, etc.) that might be apt for addressing the former issues, are not identical to the sorts of solutions (Buchanian, Ostromian, Tabarrokian,⁹ or—far more commonly, it seems—generically *statist*) that might be apt for addressing the latter. With respect to externalities, it’s the

with “any good provided by a public (i.e., state) entity.” Analysis of this “lay dogma” on public goods lies beyond the scope of this paper, however.

⁸ Elsewhere in this volume, Paniagua and Rayamajhee similarly emphasize the importance of not eliding the distinction between goods and services, on the one hand, and the impacts of their production and exchange, on the other—i.e., the importance of distinguishing analyses of the goods that generate externalities from analyses of the externalities themselves. As they state in the fifth paragraph of their [2023], one of their “points of departure” is that they “focus on externalities themselves rather than on associated goods or services.”

⁹ Alex Tabarrok [1998] has proposed the use of *dominant assurance contracts* as one potential market mechanism for surmounting some of the market failures that preclude public good provision.

actual spillover effects of actually-produced and -exchanged goods that we principally need to worry about: we wish to generate market incentives and/or state interventions to ensure that positive externalities are not too under-produced, and that negative externalities are not too over-produced. Whereas with public goods, it's the *potential* (but non-extant) non-excludable goods that we need to worry about: we wish to generate market incentives and/or state interventions that ensure that welfare-enhancing non-excludable goods still get produced, even notwithstanding the sorts of collective action problems that might otherwise inhibit private actors from spontaneously providing them at welfare-maximizing levels, in response to ordinary market incentives.

So far, so Clerical View. In the paper's remainder, I will analyze each of the previously-identified lay dogmas with reference to this framework. Where relevant, I include diagnoses of how the dogma might naturally arise out of an incautious deployment, or an imprecise grasp, of certain features of the Clerical View. Further, I catalog the risks posed by each dogma's acceptance, and illustrate the advantages that come from a firmer adherence to the central tenets of the Clerical View.

2. First dogma: coextensivism

The attentive student of political economy has likely encountered the *coextensivist* view of public goods and externalities. Though one can find evidence of coextensivism in a number of authors, the clearest expression of the view of which I am aware comes from David Friedman:

“[E]xternalities’ and ‘public goods’ are really [just] different ways of describing the same problems. A positive externality is a public good; a negative externality is a ‘negative’ public good and refraining from producing it is a positive public good. In some cases, it may be easier to look at the problem one way, in some cases the other—but it is the same problem.”

Friedman [1990, chapter 18]¹⁰

¹⁰ Other formulations of this dogma might be cited as well, but Friedman's is worth highlighting because it is especially clear and forceful in this regard. For good measure, though, I will share here another reference which, though perhaps expressing the position less directly, bears a slightly more “canonical” status. Cornes and Sandler [1996] describe public goods as a “special case” of externalities in their Preface (at page xvii; cf. also the definition offered on p. 6). Shortly thereafter—and also on p. 6—they clarify that externalities and public goods are both instances of the more general phenomenon of (market-failure-generating) “incentive structures.”

One also often encounters this view in the applied economics literature. A couple recent examples drawn from the *Journal of the Association of Environmental and Resource Economists* will illustrate. Costello and Kotchen [2022, 952] clearly betray their coextensivist commitments when they write that “Our starting point is one where an environmental externality exists (creating an environmental public bad) ...” Meanwhile, Jacobsen *et. al.* [2017] title their paper “Public Policy and the Private Provision of Public Goods Under Heterogeneous Preferences,” but it seems clear from the examples they employ that the “public goods” whose private provision they are investigating are more aptly characterized as garden-variety external effects. Their central example, in fact, centers on *literal* gardens. (This is especially literal for speakers of British English, for whom “garden” is sometimes used to mean roughly what Americans mean by “yard” or “lawn.”) On p. 245 we read that “landscaping around private homes” is a “local public good.” (It is, in fact, an instance of what they label “individually-provided public goods.”) And later they write (on p. 261) that “Privately provided neighborhood amenities such as well-kept lawns, sidewalks, and home exteriors are local public goods to which the policy implications of our model may be applied directly.” (A few sentences later they further elaborate their analysis by noting that “manicuring and fostering an extensive garden” is a more costly way of providing this public good, relative to merely “mowing the lawn and trimming.”) Furthermore, it seems clear that they are not modeling such behaviors as individual *contributions* to a broader public good, such as “neighborhood aesthetics”; rather, the outcomes of such homecare initiatives are themselves the public goods. Later, in the context of analyzing individual households' decisions regarding energy efficiency, they invite readers to “Imagine a public good like the one coming from household energy efficiency

Coextensivism, however, is an error. This is the case, first and foremost, because externalities are not “goods,” nor of course are they services. They are the external *effects* of goods and services, and it is simple sloppiness to conflate goods and services themselves with the third-party impacts generated by their production, exchange, and consumption.¹¹ If we are to be charitable, we will concede that there are instances when it seems natural to conceptualize an external effect as constituting a good in its *own* right—to look at matters in the way that Friedman asserts they can always be looked at. This charitable interpretation will be expounded (and, later, critiqued) more fully below. But for now let us simply grant that there are cases where externalities might be defensibly characterized as public goods in their own right. It still remains the case that (1) other public goods are *not* profitably modeled as externalities, and (2) many externalities cannot be modeled as public goods. Coextensivism disposes us to miss phenomena that fall into categories (1) and (2), and forces upon us the Procrustean distortion of subsuming all public goods *and* all externalities into (3) the category of externalities that *can* be modeled as public goods. Further, insofar as coextensivism leads us to overlook goods of type (2) in particular, we are apt to miss the fact that externalities—far from being non-excludable epiphenomena that we must, with resignation, regard as being enjoyed by (or, in the case of negative externalities, *inflicted upon*) everyone alike—are actually “packageable.”¹² When we recognize externalities as packageable, we are alert to opportunities whereby we might transfer them to the highest-value consumer, or to the lowest-cost sufferer.

2.1. *Diagnostics: steelmanning coextensivism*

If, as I’ve alleged, coextensivism diverges so sharply from the Clerical View and proceeds from a basic “category error”—that of conflating goods and services with their third-party impacts—how and why did it arise in the first place? In this section I supply an “error theory” that seeks to explain how the view has attained such prominence in certain quarters, notwithstanding its erroneous nature. This error theory involves an examination of the strongest and most charitable rendering of the view: a “steelmanning” of the coextensivist’s position, rather than a strawmanning.

Coextensivism can seem a natural view, I think, because of a structural resemblance that sometimes arises across the two domains of economic theorizing described in Section 1. A special case of externalities—what I shall henceforth term the “limit case”—occurs with what above we labeled ubiquitous externalities: external effects that impact every single member of the relevantly-specified community. (I borrow the term “limit case” from University of Toronto political theorist Joseph Heath, who writes in an unpublished manuscript titled “Anodyne Privatization” the following line: “A non-excludable good, for instance, is just the **limit case** of a good that generates positive externalities.”¹³) *Ubiquity* as it pertains to externalities bears a sort of irresistible structural similarity

and the associated reduction in climate change impacts for the globe” (p. 277). Again: that household’s increased energy efficiency is *itself* regarded here as the public good.

¹¹ See, again, note 8 above, citing Paniagua and Rayamajhee’s insistence on drawing this same distinction.

¹² For more on the “packageability” of public goods and services, see Rayamajhee and Paniagua [2021], especially section 3.3.

¹³ Heath [N.D., 7], **emphasis** my own. Other formulations of this claim might be cited as well, but Heath’s is worth highlighting because it is especially clear with regard to the fact that a public good is, in an interesting sense, a “limiting case” of an externality-generating good. For good measure, though, I will share here another reference which, though perhaps expressing the position less directly, bears a slightly more canonical status. Section 3.5 of Cornes and Sandler [1996] comprises a taxonomy of “special types of externalities”; they write there that “The special case [of an externality] that we call the ‘standard’ pure public good is perhaps the most common one encountered in the literature” [1996, 53].

to *non-excludability* as it pertains to goods and services. Indeed, in some cases, it might seem tempting, or perhaps even theoretically fruitful, to model a ubiquitous externality as a distinct public good: “distinct,” that is, from the good or service which itself generates the externality.

A few simple examples illustrate the naturalness of modeling ubiquitous externalities as non-excludable goods in certain cases. Consider first a wealthy citizen who, annoyed by a nearby factory’s pollution, purchases air scrubbers for it to install on its smokestacks. As a result, the air quality in the surrounding neighborhood (including, but not limited to, the area around the wealthy citizen’s house) improves dramatically. We have here a case where the citizen and the factory owner engaged in free exchange for mutual advantage—but also where, as a side-effect of their transaction, a positive externality is generated. And since this externality—clean, breathable air—is (at least for members of the relevantly-specified community) non-excludable and non-rivalrous, this positive externality just happens to be a public good as well. Consider also the following illustration: subject to some complications to be discussed more fully below, *herd immunity* might emerge as the (cumulative effect of) the external effects of a subset of a population’s vaccine consumption. When this subset reaches a certain critical mass, we might say that their shared consumption of a vaccine (itself a private good, distinct from herd immunity) generates a positive externality in the form of herd immunity. But precisely because this external effect is a ubiquitous externality—it benefits every single member of the relevant community, regardless of their individual inoculation status—it becomes possible, and tempting, to conceptualize it as a non-excludable good. Which is to say: to model it as a public good.

It may help to employ a simple mathematical model when analyzing these (and other) cases. Let our model have the following elements: a population P of n individuals $x_1, x_2, x_3 \dots x_n$, each of whom has a utility function $U(x_i)$, engaging in various forms of economic exchange E^1, E^2, E^3, \dots etc. It may sometimes transpire that parties $x_1, x_2, x_3 \dots x_i$ engage in economic exchange E^1 that has effects that enter (either positively or negatively) into the utility function(s) of external parties $x_{i+1}, x_{i+2}, x_{i+3} \dots x_{n-j}$. When this transpires, we observe (positive and/or negative) external effects of the E^1 to which individuals $x_1, x_2, x_3 \dots x_i$ are the internal parties.

The number of parties impacted by these external effects may vary, depending on the nature and location of the economic activity in question. It can range from a lone individual’s being affected (just party x_{i+1} , e.g.), to a pair or triplet of impacted individuals (parties $x_{i+1}, x_{i+2}, x_{i+3}$, perhaps) ... to a tremendous number of individuals being impacted. And, in the limit—and this is Heath’s point—the impacts might enter into the utility functions of every single one of the n members of the relevant community (however specified). In terms of the above model, this is the case where “ $j=0$.” In this limit case, there are zero members of our population who are spared from the (internal or external) effects of E^1 ... or, as we might put it, there is no one who is *excluded* from being impacted by said economic activity.

In this limit case, it might seem natural to identify the “non-exclud-*ed*” nature of the impact of E^1 with the “non-exclud-*ability*” of various goods and services: the notion that so animates the thought of public goods theorists working in the Samuelsonian tradition.¹⁴ Call this conflation—the conflation of (i) the *ubiquity* of the impacts observed in the limit case with (ii) the notion of nonexcludability countenanced by Samuelsonian public goods theory—the “General Coextensivist Maneuver.” As we have just seen, this Maneuver is sometimes defensible, as the cases of the smokestack scrubbers and herd immunity illustrated. But when the Maneuver is pressed too far, and

¹⁴ Samuelson [1954] is widely regarded as the origin of the contemporary literature on public goods.

is seen to encompass *all* instances of public goods and *any* instance of an externality, the dogma of coextensivism results. Coextensivism leaves out of account public goods that cannot be fruitfully identified with external effects of other goods. It further leaves out of account externalities that cannot be fruitfully modeled as non-excludable goods in their own right. And finally, even in some of the cases where it *does* seem defensible, like the vaccines/herd immunity case, it fails to fully capture the true nature of the relationship between external effects and public goods. The next subsection explains each of these shortcomings of the General Coextensivist Maneuver in more detail, taking them in the order just presented.

2.2. *The first dogma's three mistakes*

2.2.1. *Public goods not involving externalities*

To explore coextensivism's first infelicity—its seeming inability to recognize the existence of public goods that are not themselves fruitfully modeled as external effects of other transactions—let us imagine a small village with a mosquito problem. This village comprises precisely one hundred households occupying equally sized plots distributed evenly across the village's footprint. Consider first the following solution to the infestation: a nearby owner of a crop-duster offers to perform a summer's worth of pesticide-delivering flyovers, sufficient to keep the mosquito population at bay for a year. The price tag is \$10,000. Via the device of a dominant assurance contract, a subset of the village's households—say, 60 of them—contract to jointly produce the \$10,000 annual payment—with each household pre-pledging a contribution in an amount such that the (prospective) benefit of mosquito abatement would exceed the cost of its contribution. Thus transpires a summer's worth of welfare-enhancing, mosquito-eliminating occasional flyovers. Does the local pilot's service constitute a public good for this community? It certainly seems to: it's in the nature of this delivery technology, we'll assume, that the aerially-administered pesticides cannot discriminate between the 60 contracting households, and the 40 “free-riding households”; the service is in this respect wholly non-excludable. Does this service generate externalities? In this case, the answer to this question also appears to be “yes.” Forty households enjoy the beneficial spillover effects—the *positive externalities* (assuming for now they each regard the resulting mosquito abatement as a benefit)—generated by the transaction to which the other sixty households are internal parties. (In terms of the simple mathematical model developed above, we have here an instance where “ $n=100$,” “ $i=60$,” and “ $j=0$ ”: an instance of the Heathian limit case.)

Consider next the following alternative solution to the infestation: the village's mayor levies an annual tax of \$100 on each household to finance this endeavor. Jointly, the village's 100 households contribute \$10,000 to the cause, and thus transpires a summer's worth of welfare-enhancing, mosquito-eliminating occasional flyovers. Does the local pilot's service constitute a public good for this community? It certainly seems to, for the same reasons set forth in the preceding paragraph. But does this service involve any externalities? Plausibly, it does *not*, because the reasoning set forth just above does not apply in this case: there simply are no parties who stand “external” to this transaction, in the relevant sense. *Every* household contributes to the mosquito abatement's provision, so *no* household experiences a spillover effect. In terms of our simple mathematical model, we have here an instance where n and i *both* equal 100. The service of mosquito abatement is a public good, but it does not generate an externality, for the simple reason that there are no parties external to this transaction. As we shall see in section 2.3 below, this distinction is important, insofar as *one* of these approaches to public goods provision, but not the other, offers opportunities for creative or

innovative arrangements of a broadly Coasean nature. If we conflate effects that are non-excluded with those that are un-exclud-able, though, we are apt to overlook such opportunities.

2.2.2. *Externalities not constituting public goods*

To explore coextensivism's second infelicity—its seemingly inability to recognize the existence of externalities that cannot themselves be fruitfully modeled as further constituting nonexcludable goods in their own right—let us consider Alex, Brooke, and Charlie: all of whom live at the end of a private, gravel road. Alex and Brooke do not mind that it's gravel—though nor would they mind if it were paved. But Charlie—at the road's very end—is very much in favor of having a paved road. So s/he decides it is worthwhile to have someone pave it ... and Alex just happens to be a professional paver. Thus, Charlie pays Alex to pave the road: a paradigm case of free exchange for mutual benefit. In this event, Brooke gets a windfall gain from this mutually-beneficial transaction between Alex and Charlie: a paved road, for free. But insofar as this is a private, gated road, only these three parties enjoy it: all other members of the relevant public can be, and are, excluded from enjoying this newly-paved stretch of road.¹⁵

At this point, the coextensivist may respond as follows: “It seems you are quite mistaken to think that there are members of the ‘public’ who are excluded from enjoying this road. For as you’ve set matters up, the relevant ‘public’ here comprises only three members: Alex, Brooke, and Charlie. Expressed in terms of our preceding model, $n=3$ here: our population P has only three members, and this just is what before you called the ‘limit case’.” And in fact, the committed coextensivist might go even further, arguing as follows: “Actually, this is precisely the maneuver that we ought to make in *any* such case. We can, and should, simply stipulate, in any observed case of an externality-generating activity E , that the relevant ‘public’ be delineated precisely so as to include all and only those impacted external parties (along with all ‘internal parties’). This is the only relevant sense of ‘public’ that theorists should ever employ; certainly, it is the sense of ‘public’ that coextensivists presume when they identify ‘public’ goods with externalities.” Call this stipulative delineation of the relevant “public” the “Special Coextensivist Maneuver.”

That this Maneuver is problematically *ad hoc* can be seen from considering a slightly modified version of the preceding story. Let us now imagine that the antecedent condition of their gravel road was *ungated*, though it always sat on private land. (Though heretofore unexcluded, this road was—both in terms of technological feasibility *and* legal permissibility—manifestly excludable. It's just that its joint owners previously had never bothered to erect an exclusion device.) And let us furthermore imagine that, though it only serves to access these three homes, this particular stretch of the road also affords quite scenic views. Accordingly, it would now and then attract the random joy-riding scene-seeker, who would traverse the road for purposes other than accessing any of its three homes. Owing to its gravelly nature, however, the volume of this recreational traffic has heretofore been manageably low. The anticipated impacts of Alex's imminent paving job on the degree of “public consumption” of this road, though, have its three owner-residents worried. Now that it's to be made smooth, they fear that too many joy-riders might utilize its services, and increase public consumption to an *unmanageably high* level. Just to be safe, the three residents decide to pre-empt this possibility by

¹⁵ This example involves a positive externality, but of course the point could also be illustrated by cases involving negative externalities. The real-world case of brush fires on private lands, caused by errant sparks from locomotive steam engines—something which clearly represents a negative externality visited upon the relevant landowners, but which is in no wise a public good (or, as the case may be, “public bad”)—has been made famous in the economics literature via its extensive analysis in, e.g., Pigou [1920] and Coase [1960].

jointly hiring a contractor to erect a gate. What had previously been a (*de facto*) public good will become, by virtue of this sudden addition of an exclusion device, a private good. (Or, strictly speaking, a *club good*, as per Buchanan [1965].)

Consider now the Special Coextensivist Maneuver in the light of this example. Are we to believe that—whereas previously the “public” for the good of this scenic privately-owned gravel road comprised all members of *P* who might, at any time, have at least *contemplated* driving on it—now (post-gate-erection) said “public” comprises only the members of the *P'* whose *n* is 3: Alex, Brooke, and Charlie? How could it be that when we erect a gate, we effect such a profound shift in what constitutes the relevant public? Note further that, on the logic of the Special Maneuver, it’s actually *impossible* to “privatize” or “de-publicize” a public good. Rather, all one is doing when installing an exclusion device on a previously-non-excluded good (or whenever one encloses the Commons), is diminishing the scope of the relevant public for which this was/is a public good. (Consider: “This bridge had formerly been freely accessible. But by installing a toll booth on either end, we haven’t converted it to a private good, or to a club good. Instead, we’ve simply converted the relevant ‘public’ for whom this was a non-excluded good from ‘anyone and everyone’ to ‘anyone willing and able to pay the toll.’”) I take this result to serve as a *reductio ad absurdum* on the Special Coextensivist Maneuver. There are externalities other than ubiquitous externalities—local externalities every bit as much as there are global ones. The truth of coextensivism cannot be established by definitional fiat—by simply stipulating that the extension of the modifier “public” should always coincide with the set of impacted external parties. For if it were, it would not even make sense to speak of a given externality impacting “more” or “fewer” members of a given community—any externality would necessarily (by definition) affect *every* member of the (relevantly-delineated) public. This Maneuver can establish the truth of coextensivism only at the cost of radically distorting many other fixed points in our common ways of talking and thinking about externalities.

2.2.3. *Obscuring the existence of “publicized” (or “impure,” or “co-produced”) goods*

Coextensivism’s third infelicity arises in cases wherein the General Coextensivist Maneuver might initially appear to be legitimate, but where closer analysis reveals that even here an important conflation has occurred. In particular, this conflation can arise in cases of what various writers have variously termed “impure” public goods, “co-produced” public goods, or (as I myself prefer to describe them) *publicized* goods.

In the words of Vaughn Bryan Baltzly [2021, 376], a publicized good is “any whose ‘public’ character results only from a policy decision to make some (otherwise private) good freely and universally available, rather than from any features that are (we might say) *inherent in* or *intrinsic to* the good itself.” We might conceive them as otherwise-perfectly-excludable goods *g* that are “publicized” via the device of appending, to their names and/or descriptions, prefatory phrases like “guaranteed universal access to” *g*, or in some instances, “induced universal (or at least *widespread*) consumption of” *g*. As an illustration of the former, Baltzly [2021, 383-4] notes how the social safety net can be modeled as a publicized good generated by the state’s guarantee of universal access to certain otherwise purely excludable, paradigmatically private goods: in this case, certain forms of *insurance* coverage (unemployment benefits, disability- and old-age-insurance, etc.). And as an illustration of the latter, Baltzly [2021, 385-6] notes how herd immunity can be modeled as a publicized good generated by community members’ widespread consumption of vaccines. Adapting Baltzly’s formalism just slightly here, we may denominate such publicized goods with a capital Greek letter, and say that a publicized good like herd immunity is the Γ emerging from the external effects of

many individuals' consumption of the g of vaccines. Getting a shot of, e.g., a COVID vaccine constitutes the consumption of a private good. But we can publicize this good by encouraging, or even mandating, universal (or at least widespread) consumption of this private good, in pursuit of some other (analytically distinct, yet still closely related) public good. In this case, herd immunity is the non-excludable good that results from the aggregation of all the external effects of many individuals' vaccinations.

It is therefore tempting to make the General Coextensivist Maneuver here, and to conceptualize herd immunity as being itself an externality created by vaccinations. But even in this instance, the Maneuver is technically a mistake: there is a subtle distinction between *the external effects of individual vaccinations* and *the public good of herd immunity*, and the Maneuver elides this distinction. What in fact transpires in this case is better described as follows: A vaccine jab is a perfectly excludable good that creates its own wholly private benefit: personal inoculation against a disease like COVID-19. But the wholly private benefit of inoculation also has a spillover effect that we can easily and naturally characterize as a positive externality—"self-elimination-as-a-potential-disease-vector," we might call it—insofar as now, once inoculated, one won't transmit that disease to any close contacts. So your close contacts enjoy the positive externality generated by your personal disease immunity. But furthermore—though of course only above a certain threshold—the accumulation of all these positive external effects might eventually generate a distinct, nonexcludable public good: herd immunity. Note that, unlike the "local" externality of vector-elimination—which benefits only those who might otherwise risk catching the disease from you (your close contacts, in other words)—herd immunity operates as a sort of "global" (ubiquitous) externality—it benefits *every* member of the community (here tacitly defined as the relevant "herd"). In terms of the simple formal model developed above, we would say that the benefits of herd immunity extend to every single member of the population P . Insofar as your own personal inoculation is partially constitutive of the effective elimination of P as a reservoir of the disease, you play a role in benefiting everyone who stands to gain from said elimination—even those who might *not* be potential close contacts, and who were thus unlikely to capture any of the "localized" externalities generated by your individual inoculation.

Besides its already-noted affinities with the framework set forth in Baltzly [2021], this example is also quite similar to what Cornes and Sandler [1996, 57] label an "impure public good": "an individual commodity that appears twice in the consumer's utility function, once on its own as a private commodity, and once in combination with the quantities consumed by others, thereby forming a public good or bad." (Neither the brief introduction to the notion at their [1996, 57-8], nor their more technical analysis at pp. 255-72, seems to fully distinguish what here I have labeled as "local" versus "global" externalities, however.) This example is also akin to an instance of what the Bloomington School theorists call "co-produced" public goods. (For more on co-production, see (e.g.) Paniagua and Rayamajhee [2021] section 3, and Rayamajhee and Paniagua [2021] section 3.2.) Whether we conceptualize these phenomena as publicized goods, or as impure public goods, or as co-produced public goods, coextensivism disposes us to overlook the fact that very often "spillover effects of $g \neq \Gamma$ ": that is, that the externalities generated by any one individual's consumption of any single unit of some g , is analytically distinct from the (impure) public(ized) good that might get (co-)produced in the process.

To see that the foregoing analysis does not serve merely to capture a "distinction without a difference," simply note that very different policy instruments will suggest themselves, depending on whether we conceptualize the problem at hand as one of internalizing externalities, or one of public goods provision. In the former case, we might worry that the (local) external benefit associated with

disease inoculation—vector-elimination—will be under-produced in a purely free-market setting: too many individuals will regard the marginal benefit of a vaccine to be exceeded by its marginal cost, and will therefore decline to get vaccinated. But the proper set of broadly Pigouvian interventions—subsidies to reduce (or eliminate) the vaccine’s monetary costs; ample delivery-sites to reduce (or eliminate) various transaction costs (the inconvenience associated with finding and transporting oneself to a vaccination station, and with waiting in a queue, etc.)—will increase the number of individuals for whom now $MB > MC$, and so more positive local spillover effect will be created. In the latter case, on the other hand: if what we’re really aiming for is herd immunity, or the elimination of P as a reservoir for a given disease, then such Pigouvian tools may not be adequate, as they may not themselves get us up to an adequate threshold of inoculated individuals, sufficient to produce the relevant publicized good. In that instance, public authorities might need to resort to more drastic policy levers: vaccine *mandates*, for example.

We can see this same dynamic at play in cases of negative externalities as well: the pollution generated by some otherwise economically-productive manufacturing process (to recycle the example employed above) provides an illustration. The byproducts and outputs of a manufacturing facility—smoke, e.g.—may have negative impacts on those living in the factory’s immediate environs, via effects like noxious aromas, dirty and less-breathable air, and occluded sunlight. These are the local externalities. But these same pollutants might also be partially constitutive of some wider phenomenon that results from the accumulation of many such inputs—say, the depletion of the ozone layer, or the generation of acid rain. These are the global externalities. To see that these two impacts are indeed analytically distinct (and furthermore that they are worth distinguishing), it’s helpful to note what occurs in the relevant counterfactuals. In the counterfactual world in which this local manufactory is the *only* enterprise producing such pollutants, the localized effects would remain as relevant and impactful as they are in the actual world: the local air would be just as dirty for the plant’s neighbors. But the global externalities would be non-existent: absent a critical mass of co-conspirators, the pollutants this plant puts up into the earth’s atmosphere simply don’t come close enough to the relevant threshold at which the cumulative effects start to emerge. (One plant’s emitting some CFCs is not sufficient to pierce a hole in the ozone layer, e.g.) Conversely, in another counterfactual world where we remove just that one plant, its immediate neighbors are instantly relieved of the negative local externalities that they suffer in the actual world. But they still suffer the same negative global externalities. For if all those *other* relevantly similar plants, all across the globe, continue spewing forth *their* pollutants, those neighbors will still suffer the same (or very nearly the same) degree of impact from, e.g., acid rain.¹⁶

So that is the last of the first dogma’s three major mistakes: the General Coextensivist Maneuver sometimes represents a distortion, even in its best use case—the case of (what Baltzly calls) publicized

¹⁶ Obviously, the local/global distinction I’ve been deploying here is a matter of degree, not a firm distinction. The spillover effects from one party’s consumption of a private good can impact just one other party (a purely local externality), or they can impact every single member of the relevant community (a wholly global externality), and of course anything in between. This continuum is nicely captured by the scalar nature of the horizontal dimension of Figure 1 in Paniagua and Rayamajhee [2023]. (This is what they term the “scale/size of externalities” dimension. While their Figure 1 depicts this variable as discrete—cleaving into the categories of “small-”, “medium-”, or “large-scale” externalities—it is clear from their main-body text (e.g., Section 2 of their paper) that this tripartite division is simply an artifact of their diagram, and that really they regard this variable as continuous.) We can also express this observation in terms of my formal model from above. Where the spillover impacts only party “ $i+1$ ”, we have the limit case of a (very) local externality. When the spillover impacts every single other member of the relevant public—the case where $j=0$ —then we have the limit case of a wholly global externality. In reality, of course, most cases will fall somewhere between these polar extremes, and spillover effects will be more-or-less local or more-or-less global, to varying degrees.

goods. (Or what Cornes and Sandler [1996] call “impure public goods,” or what the Bloomington School calls “co-produced public goods.”) The first two mistakes resulted from the fact that Maneuver is wholly inapt—and the resulting coextensivist framework entirely Procrustean—in the cases of externalities that don’t themselves seem to constitute distinct non-excludable goods, and in the cases of public goods that *cannot* be fruitfully modeled as the external effects of economic exchanges to which other parties are internal.

2.3. *(Economic) externalities without the (first) dogma: “Coasean clarity”*

One benefit to be gained from my analyses is increased perspicuity regarding the sorts of phenomena described in section 2.2. But there are other benefits at stake too—including some that involve preempting potentially costly practical errors. So I now wish to argue that the rejection of coextensivism is of considerable theoretical utility in an additional way, insofar as it affords us some degree of “Coasean clarity.” That is, it provides us an awareness of the possibility of package-able externalities, and of the Coasean solutions to which such an awareness might give rise. Another way of putting this is that, while in the grips of the coextensivist dogma, we are apt to “under-Coase.” That is, we might be insufficiently attentive to the possibility of identifying mutually beneficial negotiated resolutions of problems associated with externalities—resolutions of the sort famously associated with the paradigm inaugurated by Ronald Coase [1960].

As previously stated: insofar as coextensivism leads us to overlook the sorts of distinctions discussed above, we are apt to miss the fact that some classes of externalities—far from being non-excludable epiphenomena that we must, with resignation, regard as being experienced by everyone alike—are actually packageable. When we recognize externalities as packageable, we are alert to opportunities whereby we might direct them towards their highest-value consumer, or to their lowest-cost bearer. To see this most clearly, let us return to our discussion in 2.2.1, where we surveyed the different ways of conceptualizing—and, importantly (and relatedly), the resultant different ways of *provisioning*—the service of mosquito abatement. The discussion there might also have invited the charge that I was trafficking in “distinctions *sans* difference.” Admittedly, this charge might have seemed fairly warranted at the time. But there *is* I believe a case to be made that the distinction-mongering is merited.

To illustrate: nearly 50 years ago, in the pages of this same journal, Kenneth Goldin analyzed this very same service (insect abatement), showing that a creative and open-minded approach thereto might reveal alternative means of delivering the good or service in question, in ways that altered its character as a public good, or modified its “externality profile.” For example: let’s imagine for the moment that some households in the village *don’t* regard the aerial delivery of pesticides as a benefit. (Perhaps they *like* mosquitos. Or—and rather more plausibly—perhaps they simply wish to avoid exposure to unnecessary chemicals, and/or prefer not to kill the mosquitos (even if they might still wish to repel them).) For their sakes, we might consider adopting Goldin’s suggested solution of plot-by-plot pesticide application, utilizing the targeted device of *spray guns*, rather than coarse-grained delivery devices like crop-dusters. True, even the relatively fine-grained contrivance of plot-specific application of pesticides might still generate externalities: if my neighbors each have the chemicals applied to their plots, I myself will likely notice a substantial decrease in my local mosquito population (and a non-negligible increase in my own exposure to unwelcome chemicals). And perhaps if a critical mass of households opts for such individualized applications—a sufficient number of households, sufficiently widely geographically distributed across the entire village (as opposed to narrowly concentrated in a single neighborhood)—we may wind up effecting “universal

mosquito elimination” anyway. (This would result via a dynamic that’s very much like the one effecting herd immunity when a sufficiently large subset of a population becomes inoculated against a disease.) By attending carefully to the distinction here—by refusing to conflate public goods and externalities—we can see how a welfare-enhancing compromise might be brokered: the families most averse to experiencing even the spillover effects of their neighbors’ localized pesticide applications *could* perhaps exchange plots of lands, and cluster together in a portion of the village that would be relatively distant from their pesticide-loving fellow citizens. This arrangement would simultaneously generate both (i) a more efficient leveraging of the spillover effects of individual-plot sprayings (because all such plots would be physically adjacent), and (ii) a decrease in negative externalities, insofar as all the households who wished to avoid the spillover effects of their neighbors’ pesticide usage would be much more sheltered from these external effects.

It’s worth noting that the notion of packaging negative externalities in such a way that they might be borne by the lowest-cost sufferer is nowadays tainted, at least in the popular imagination,¹⁷ by its association with the infamous “Summers Memo.” (In a 1991 memo signed by Larry Summers (though written by Lant Pritchett), then chief economist at the World Bank, it is suggested that perhaps high-polluting industries might best be exported to the developing world, insofar as—owing to considerations including their lower levels of general economic activity, and the shorter lifespans of their inhabitants—the polluting activities impose fewer external costs in these locales, relative to the costs they would impose on inhabitants of the developed world. Understandably, large portions of the commentariat were outraged when the contents of this memo were revealed.¹⁸) Nevertheless, I believe that this notion of packaging externalities, and searching for their highest-value/lowest-cost beneficiaries/sufferers, merits more consideration than it standardly receives. The reluctance to do so, however understandable, may lead to our missing out on a whole range of creative, enterprising, innovative solutions, of a broadly Coasean nature, to a wide range of problems concerning externalities. Besides the hypothetical example related to mosquito abatement discussed just above, the example of siting apiaries next to orchards is a familiar one from the literature.¹⁹ And there are other real-world examples here too—even some that we might characterize as “Summers Memo-adjacent.” For example, some countries export some of their nuclear waste, on grounds that the recipient regions represent the “lower-cost sufferers,” owing to the fact that they are less densely populated than is any part of the exporting nation.²⁰ (Here the packaging and exporting of the negative externality strikes many as less problematic than that contemplated in the Summers Memo, insofar as the principal condition that makes the “waste-importing regions” more attractive destinations for storing nuclear waste—*viz.*, their lower population density—seemingly does not betoken any antecedent injustice or power disparity between exporter and importer.) And at least one recent study documents real-world instances of Coasean bargaining being used to fruitfully govern complex externalities in the sphere of environmental impact: the case studies found in

¹⁷ Or perhaps “popular *academic* imagination”: it’s not clear how widely-known is this memo, outside of the economically-informed intelligentsia.

¹⁸ A Wikipedia entry supplies an overview of this affair, and includes the text of the memo itself: https://en.wikipedia.org/wiki/Summers_memo. For their part, Summers and Pritchett maintain that the memo was in crucial respects intended as sarcasm or parody; Summers’ side of the story is succinctly presented in a contemporaneous article in the Harvard *Crimson*: <https://www.thecrimson.com/article/1992/2/25/summers-memo-ironic-pa-harvard-professor/>.

¹⁹ Indeed, as Steven Cheung [1973] has taught us, orchard owners and beekeepers in Washington state (and presumably elsewhere) *have* learned to package, and to trade, the (beneficial) external effects of their commercial enterprises: orchard owners contract to have beekeepers transport their hives into their orchards for purposes of promoting pollination. (Obviously, the beekeepers benefit likewise, from their bees’ access to nectar.)

²⁰ I am indebted to a reviewer from *Public Choice* for making me aware of this example.

Deryugina et al. [2021] include polluters purchasing nearby lands, payments for ecosystem services, and land acquisitions to protect the supply of drinking water.

3. Second dogma: externality profligacy

Coextensivism about public goods and externalities is often conjoined with a potentially even more insidious dogma: the *profligate* identification of externalities, or “externality profligacy.” Political economists are problematically profligate when they see so-called externalities generated by just any activity whatsoever that creates un-borne or un-compensated costs or benefits for someone other than the actor. Though one can find evidence of this tendency in multiple authors, the clearest example of which I am aware comes from Jonathan Gruber:

“An externality occurs whenever the actions of one party make another party worse or better off, yet the first party neither bears the costs nor receives the benefits of doing so.”

(Gruber [2019, 120])²¹

It will immediately strike the reader that this conception of an externality is extremely broad: *anything* that you do that I don’t like now qualifies as a negative externality, and of course vice versa.

Externality profligacy, however, is an error. We should instead restrict our attention to externalities to cases where third-party costs and benefits arise from *first-/second-party* mutually-beneficial exchange. We might put this point by saying that our economic theorizing should be restricted to what we might call “economic externalities” *per se*, as opposed to what we might call “social externalities.” We can distinguish these two categories of externalities with reference to the price mechanism. While many human exchanges are negotiated and consummated without the benefit of the information and incentives embedded in the price signal, many other transactions—including most “commercial interactions” and interactions among strangers—are thusly mediated. Only the ones in the latter category generate what we should label economic externalities, properly so-called.²²

²¹ Once again, additional sources might be cited here as evidence of the pervasiveness of this dogma. But once again, I shall rest content with buttressing the case with citations from a single source—one which, whatever it might lack in forcefulness of expression, it compensates for in its “canonicity.” Cornes and Sandler [1996] first define externalities as any “uncompensated interdependencies” (p. i), then later as “the action of one economic agent [that] influences the utility or production function of another, and [for which] no mechanism for compensation exists” (p. 5). Eventually, they cite with approval a definition offered by Meade [1973]: “An external economy (diseconomy) is an event which confers an appreciable benefit (inflicts an appreciable damage) on some person or persons who were not fully consenting parties in reaching the decision or decisions which led directly or indirectly to the event in question” (cited at Cornes and Sandler [1996, 39]).

²² Though the linguistic device of “economic vs. social externalities” is not commonly employed in the literature, the distinction it marks *does* seem to be increasingly taken note of. For example, Cornes and Sandler attest to the historical reality of a sort of “conceptual inflation” with respect to externalities when they write [1996, 6] that, subsequent to Pigou, “the notion of externality encompassed an ever-increasing variety of economic situations.” And for two contemporary examples of authors who *do* use the term “social externality” in approximately this way, see Vogler [n.d.] and Fleurbaey, Kanbur, and Viney [2020]. The former argues that, at least in certain analytic contexts, it is useful to have a “definition of externalities ... broader than the traditional perspective because it is not limited to voluntary economic exchange, but also applies to many other forms of social interaction.” (Vogler then proceeds to define “economic exchange” as “all market transactions of labor, goods, services, and currencies that two or more parties voluntarily engage in and that are associated with prices,” and “social interaction” as “voluntary and involuntary encounters of two or more people in a narrow space that involve oral communication or physical contract or both.”) (All quotations taken from p. 2.) The latter

(We might also invoke in this vicinity a distinction between *market* and *non-market* exchange.) Among other things, this rules out the possibility of labeling as an externality any given aspect of any given person's individual behavior that any other given individual might find offensive or problematic. Otherwise, we run the risk that the (Pigouvian etc.) tools of economic analysis and policy-making be brought to bear in order to legitimate all *manner* of interferences with liberty—to bring state regulation into the sphere best left to manners and morals. Put another way: while in the grips of the dogma of externality profligacy, we are apt to “over-Pigou.” That is, we may display unwarranted enthusiasm for the sorts of tax- and subsidy-based resolutions of externalities-generated market failures, famously associated with the work of A.C. Pigou [1920].²³ There is a further risk that externality profligacy will lead us to “over-Coase”: to display unwarranted enthusiasm for an overly-transactional, bargaining- and negotiation-based approach to resolving disputes with our peers. The next two sub-sections explore each of these dangers more fully.

3.1. *The second dogma's first danger: over-Pigouing*

If personal liberty is to have any meaningful content, we must be at liberty, at least some of the time, to act in ways that give offense or annoyance to (at least some of) our fellows—to act in ways, that is, that create social externalities. So long as these offenses and annoyances don't transgress into the territory of *harms* (or, better yet, *wrongs*) they must be tolerated in the name of liberty. Of course, how to properly demarcate the line separating *harm* from *offense* (and *nuisance* and *inconvenience* and other related categories) is an enormously fraught issue, upon which it is basically hopeless to expect consensus. But to deploy the notion of social externality in the way that Gruber seemingly endorses, is to invite the drawing of this line in a way that errs far too much in the wrong direction. The worry is that it engenders a sort of “credibility transfer” from the garden-variety application of standard Pigouvian maneuvers. So, e.g., we might all agree that Pigouvian taxes on certain industrial activities are a justifiable way of inducing industrialists to (approximately) internalize the full social costs of their polluting (but otherwise socially-profitable) activities. Yet from the independent intuitive plausibility of *these* sorts of welfare-enhancing Pigouvian economic interventions, we might then be tempted to (illicitly) conclude that structurally-similar “externality-internalizing” interventions would be merited in the case of, e.g., my neighbor's playing her music too loudly. And I hold that we must hold the line in this regard; to do otherwise is to run the risk of over-Pigouing—of bringing the (often coercive) tools of public policy to bear on what are essentially private matters.

To illustrate, I shall elaborate upon the example, involving my own neighbor, to which I alluded just above. When “Rhonda” plays her music while she does yardwork or sits on her back porch, all the members of my household can hear it. We have a range of responses available to us when we do—from embracing the nonvoluntary soundtrack, to tolerating it when it might not be all that welcome (but is not all that much of a nuisance), to requesting that she turn it down when it's preventing us from sleeping at night. And these are the sorts of responses that it is fitting for us, as neighbors and

propose to “further enlarge the [traditional] perspective [on externalities] and start thinking about a broader framework in which *any pattern of influence of an agent or a group of agents over a third party, which is not mediated by any economic, social, or psychological mechanism guaranteeing the alignment of the marginal net private benefit with marginal net social benefit*, can be attached the ‘externality’ label and be scrutinized for the likely negative consequences that result from the divergence” (page 6, *emphasis* in the original).

²³ To say that relatively few social externalities merit resolution through standard economic policy-making tools, though, is not thereby to imply that all *economic* externalities *are* the proper subject of Pigouvian intervention. To cite just one important example, pecuniary externalities are no doubt properly classified as negative *economic* externalities, but few economists would endorse government intervention to mitigate these effects of market competition.

inhabitants of shared sonic space, to exhibit—and the sorts of responses that it is fitting for Rhonda to expect from us (and to prospectively keep in mind, on those occasions when she decides to “share” her music with us). But the move to conceive this as a “social externality”—an uncompensated benefit that Rhonda provides us on the days we’re enjoying her tunes, and an un-borne cost on those days when we’d rather not—invites us to seek measures (implemented by third parties) whereby her incentive structure is modified in such manner that she “internalizes” these costs (and benefits). But were such a comprehensive program of third-party-mediated incentive-altering policy-intervention to be successfully carried out, Rhonda would thereby cease to be our *neighbor*. All things considered, I hold, it is better to leave these matters to the sphere of conversation, persuasion, and shared accountability. To put the point a bit too crudely—though perhaps most clearly—the insight here is much the same as that which underlies our tendency to counsel children, once they reach a certain age, not to be too much of a “tattle-tale”: “You know,” we might tell them, “it’s sometimes better to handle these sorts of inter-personal differences and conflicts directly, rather than always running to a teacher or a parent for intervention.”

3.2. *The second dogma’s second danger: over-Coaseing*

At this point, a natural response to this worry might be to extol “Coase-ing” as an alternative to thusly (over-)Pigouving. (What could be less interventionist and “tattle-tale-ish” than approaching the offending party directly, and opening up negotiations with him or her?) But I want to argue that a broadly Coasean approach is actually inadequate here as well. Specifically: there would be serious social costs to an arrangement wherein we were over-disposed to seek (Coasean) bargains with our neighbors and fellow citizens—bargains whereby we sought to mitigate the costs of others’ personal activities in the most cost-effective and mutually-acceptable manner. Indeed, the worry is that—whereas with economic externalities we risk under-Coase-ing in the manner described in section 2.3 above—with *social* externalities we risk “over-Coase-ing,” in a manner I shall now describe.

The values of neighborliness and the norms of community can be undermined by over-Coaseing as well as by over-Pigouving. In sum, the argument holds that an excess of Coasean bargaining might serve to render our relations with our neighbors unduly *transactional*. Consideration of the following hypothetical illustrates the intuition here. Imagine a software developer to have created a “Coase App” that any smartphone user can install on their mobile device. It works as follows: any time party A regards party B’s activity as a nuisance, A can use the app to anonymously send B a proposed Coasean bargain: “I wish to be compensated to the tune of $\$M$ for the unpleasantness of having to endure your music,” e.g. The app knows the value of M for A—A had to enter it before s/he could send the proposal—but does not disclose this value to B. Instead, it asks party B “How much would you be willing to pay to continue playing your music?” If the amount N that B enters here exceeds M , then the app (which as a condition of use must be linked to its user’s bank account) “splits the difference”: it announces to both parties that B may continue playing their music, but it also withdraws $\$[M + \{(N-M)/2\}]$ from B’s account and transfers it to A. Likewise, if the amount B enters is *less* than M , then the app announces that party B is the least-cost avoider, and decrees that s/he should turn off the music. But then in compensation for B’s concession, the app also transfers $\$[N + \{(M-N)/2\}]$ from A to B.

To illustrate using the case of Rhonda described above: I might use the Coase App to request from her at least $\$8$ per hour of compensation for her evident intention to spend the afternoon lounging poolside while listening to the *Phantom of the Opera* soundtrack. But in reply—should she decide to view and consider the request on the app—she might communicate, also via the app, that it’s worth

\$12/hour to her to be able to enjoy her poolside Andrew Lloyd Weber. In that case the app would declare me the lowest-cost avoider, but then would pay me \$10 (of Rhonda's money) for every hour that Rhonda plays her music. In this event, both Rhonda and I thereby enjoy a \$2 "consumer surplus" of sorts. (I am compensated \$2/hour more than I had requested, and Rhonda pays \$2/hour less than the value to her of her music.) And had Rhonda's *N* been only \$4, then the Coase app would have transferred \$6/hour from me to her, in exchange for her muting of her stereo.

Obviously, many practical difficulties would have to be surmounted for such an app to work effectively. (Among other things: some contrivance, no doubt utilizing the app's location services (which perhaps must always be enabled?), must be utilized to handle cases where party A either is ignorant of B's identity, or lacks B's contact info. Issues of enforcement and compliance must also be carefully considered. (What if B winds up accepting A's offer, but fails to cease and desist the nuisance-generating activity?) And needless to say, some care must be taken to eliminate the incentive to engage in rent-seeking, entrepreneurial "nuisance-generating" behavior (just as in other arenas, provisions are taken to forestall opportunities for extortion).)

But even assuming that such practical difficulties could be surmounted: is a world where the Coase App is ubiquitous, and where its use to effect localized Coasean "micro-bargains" is pervasive, a better world than our own present one? According to this counter-Coasean, "anti-over-transactionalizing" argument, the answer is "no." Fears that such a social world would be unduly "transactional" are not unfounded, and this argument has, at least *prima facie*, some genuine force.

3.3. *(Social) externalities without the (second) dogma: a "Coasean ideal of community"*

There is an additional argument for this same conclusion that is of a broadly "Coasean" nature, however. Unlike the considerations invoked in section 3.2 above, though, it is "Coasean" more in the sense of deriving from his classic 1937 essay "The Nature of the Firm," than from his 1960 "The Problem of Social Cost." For many of the sorts of negotiation and bargaining and compromise that we engage in when we jointly create and navigate a shared social environment with our fellows and neighbors, though they can be modeled as "transaction costs," are in fact *welcome* transaction costs. Rather than the sorts of costs that we should seek to avoid or minimize, they are, rather, the very "stuff of sociality." They are how we build and experience a shared life together. When we model these sorts of unintended byproducts of in/voluntary encounters with our fellows as social externalities, we thereby seem to invite the application of (at least the theoretical apparatus for thinking about) third-party adjudication to mutually modify our behaviors so as to "internalize these externalities."²⁴

The Coasean element (in the sense of Coase [1937]) of the present argument involves the idea of *optimal transaction costs*. Just as firms arise from the need to internalize certain transaction costs—and just as the optimal size of the healthy firm is driven by efforts to achieve the optimal balance between

²⁴ At this point, though, it is important to consider a rival view of things. For it may well be the case that a citizenry's facility with the notion of "social externalities" could lead to behavior that is more thoughtful and considerate. Pervasive recognition of the fact that personal behavior generates uncompensated external effects might manifest in a stronger desire not to (unintentionally) harm others in one's immediate environment. Widespread uptake of the notion of social externalities might thereby substantially *reduce* the need to ask for third-party enforcement. (I am indebted to an anonymous reviewer for recognizing this possibility, and in particular for this way of expressing it.) While this perspective undoubtedly has merit, I remain unpersuaded that the *net* effect of widely adopting such an economic, transactional, "market-ized" conception of all social relations would be pro-social.

“internal” and “external” transaction costs—similar remarks might be made with respect to the existence, and the optimal size of, healthy *communities*. For any community is going to comprise both (i) institutions of shared governance whereby certain transaction costs are eliminated simply by preempting certain (social-)externality-generating behaviors, *and* (ii) norms and expectations regarding direct, peer-to-peer engagement to effect resolutions of other, “lower-stakes” issues of (social) externalities. And just as it’s rarely the case that the optimal firm size is “1,” so also is it rarely the case that “all” represents the optimal amount of the social transaction costs that should be eliminated by our institutions of shared governance, or of the social externalities that these institutions should seek to internalize. Like a Coasean firm, the institutions of shared governance that constitute our communities should be “right-sized.” We want our (negative-)externality-internalizing institutions to be at efficient scale. To put the point now in terms more commonly associated with Hayek [1976] and Buchanan [1964], rather than in Coasean terms: If our externality-internalizing institutions are too limited in scope, we forfeit catallaxy, and all the benefits provided by the extended market order. (For the promotion of our welfare, material and otherwise, we would depend solely upon whatever assistance might be proffered by those with whom we have trusted personal relations. We would return, that is, to the days of the hunter-gatherer tribes.) However, if these institutions are too *great* in scope—in the sense of being too *intrusive*, such that they could be brought to bear on all manner of social minutiae, no matter how trivial—then we forfeit our sense of community. We have no *neighbors* in an overly-Coasean (and, for that matter, an overly-Pigouvian) world: instead, we have only an inexhaustible supply of potential *business partners*.²⁵

4. Summary and conclusion: the governance of (dogma-free) complex externalities

I have argued that we need to get clearer on the relation between public goods and externalities in two ways. First, I have argued that, while some externalities can be modeled as public goods and vice versa, sometimes public goods *don’t* involve externalities, and sometimes externalities *aren’t* public goods. This seems to be a fundamental point concerning the nature of externalities that has been under-appreciated by many who occupy positions in the broader political economy space, and even (surprisingly) by some contemporary working economists. And second, I have argued that—lest, via the familiar phenomenon of “concept creep,” all manner of activities which should otherwise reside in the protected sphere of personal liberty become subject to public regulation—we should remain vigilant against externality profligacy, and insist on restricting the notion to its original site of application: third-party impacts of market exchanges that are, for the first two parties, welfare-enhancing.

Failure to excise the first dogma from our thinking—failure, that is, to recognize that many externalities are entirely excludable, and thus perfectly packageable—puts us at risk of under-Coaseing, in the sense that we might neglect to direct the external effects of our economically-productive activities towards the lowest-cost avoiders and/or the highest-benefit recipients. Meanwhile, failure to excise the second dogma from our thinking—failure, that is, to recognize that not every “cost” of others’ activities properly occasions the application of (sometimes coercive) economic policy levers—puts us at risk of over-Pigouving and over-Coaseing. In a world where various dynamics (not least of which is the advent of social media) incentivize the officious and zealous discovery of the inconveniences, offenses, nuisances, and even “harms” generated by the exercise of

²⁵ Daniel Halliday and John Thrasher [2019, 249] make a similar point when they discuss the idea (which they attribute to some unpublished work of political philosopher Jonathan Wolff) that even the most enthusiastic partisans of markets are likely to wish for *some* spheres that remain insulated from market forces. They use the example of passengers on crowded, standing-room-only subway car, buying and selling seats on their daily commutes, as an illustration.

others' liberties, such an overly-profligate conception of externalities invites opening a Pandora's box of government bureaucracy and intervention, and all the subsequent public choice risks and threats that this implies. (In this vein, see Leeson and Thompson [2021]. Heath [2020, 304-25] also supplies a nice overview of this dynamic.)

As the other papers in this volume make abundantly clear, the governance of complex, genuine externalities is challenge enough. Let us not multiply our troubles, therefore, by employing oversimplified understandings in the attempt to govern simplistic, ersatz externalities.

Acknowledgements

For helpful suggestions on earlier versions of this work, I must thank audience members at the March 2022 Public Choice Society conference in Nashville (particularly Jan Vogler), and at the November 2022 meeting of the Philosophy, Politics, and Economics Society in New Orleans (particularly Nathan Goodman and Matthew Nagler). The guest editors of this special volume—Ilia Murtazashvili, Pablo Paniagua, and Veeshan Rayamajhee—each supplied feedback, not only on my conference presentations (Veeshan in Nashville, Ilia and Pablo in New Orleans), but also on earlier drafts of the paper (with Pablo having done multiple iterations). Finally, the formal review process with *Public Choice* proved enormously productive in terms of improving this work. For their role therein, I must thank the aforementioned guest editors, five anonymous reviewers, and editor-in-chief Peter Leeson.

References

- Baltzly, V. B. [2021]. Concerning publicized goods (or, the promiscuity of the public goods argument). *Economics and Philosophy* 37(3): 376-94.
- Buchanan, J. M. [1964]. What should economists do? *Southern Economic Journal* 30(3): 213-22.
- Buchanan, J. M. [1965]. An economic theory of clubs. *Economica*, 32(125): 1-14.
- Cheung, S. N. S. [1973]. The fable of the bees: an economic investigation. *The Journal of Law and Economics* 16(1): 11-33.
- Coase, R. H. [1937]. The nature of the firm. *Economica* 4(16): 386-405.
- Coase, R. H. [1960]. The problem of social cost. *Journal of Law and Economics* 3: 1-44.
- Cornes, R., and T. Sandler [1996]. *The theory of externalities, public goods, and club goods* (2nd edition). Cambridge, UK: The Cambridge University Press.
- Costello, C., & M. Kotchen [2022]. Policy instrument choice with Coasean provision of public goods. *Journal of the Association of Environmental and Resource Economists* 9(5): 947-980.
- Deryugina, T., F. More, & R.S.J. Tol [2021]. Environmental applications of the Coase Theorem. *Environmental Science and Policy* 120: 81-88.
- Fleurbaey, M., R. Kanbur, and B. Viney [2020]. Social externalities and economic analysis (August 06, 2020). The SC Johnson College of Business Applied Economics and Policy Working Paper Series No. AEP WP No. 2021-07, Available at SSRN: <https://ssrn.com/abstract=3933271> or <http://dx.doi.org/10.2139/ssrn.3933271>
- Friedman, D. [1990]. *Price theory: an intermediate text*. Cincinnati: South-Western Publishing. Also available at http://www.daviddfriedman.com/Academic/Price_Theory/PThy_ToC.html.
- Goldin, K. [1977]. Equal access versus selective access: a critique of public goods theory. *Public Choice* 29(1): 53-71.

- Gruber, J. [2019]. *Public finance and public policy*, 6th edition. New York: Worth.
- Halliday, D., and J. Thrasher [2019]. *The ethics of capitalism*. New York: Oxford University Press.
- Hayek, F. [1976]. *Law, legislation, and liberty*, Vol. 2. Chicago: Chicago University Press.
- Heath, J. [2020] *The machinery of government: public administration and the liberal state*. New York: Oxford.
- Heath, J. [N.D.] Anodyne privatization. Unpublished manuscript, presented to the PPE Society's "Ethics & Public Policy Working Group," Feb. 18, 2022: <https://ppesociety.org/epp-working-group/>.
- Huemer, M. [2013]. *The problem of political authority*. New York: Palgrave Macmillan.
- Holtermann, S. E. [1972]. Externalities and public goods. *Economica* 39(153): 78-87.
- Jacobsen, M., J. LaRiviere, & M. Price [2017]. Public policy and the private provision of public goods under heterogeneous preferences. *Journal of the Association of Environmental and Resource Economists* 4(1): 243-280.
- Leeson, P. T., & Thompson, H. A. [2021]. Public choice and public health. *Public Choice*, 1-37.
- Mankiw, N. G. [1998]. *Principles of economics*. Fort Worth: Harcourt Brace College Publishers.
- Meade, J. E. [1973]. *The theory of economic externalities. The control of environmental pollution and similar social costs*. Geneva: Sijthoff-Leiden.
- Paniagua, P., & V. Rayamajhee [2021]. A polycentric approach for pandemic governance: nested externalities and coproduction challenges. *Journal of Institutional Economics*, 1-16.
- Paniagua, P., & V. Rayamajhee [2023]. On the nature and structure of externalities. *Public Choice*. <https://doi.org/10.1007/s11127-023-01098-1>.
- Pigou, A. C. [1920]. *The economics of welfare*. London: Macmillan and Co., Limited.
- Quine, W. V. O. [1951]. Two dogmas of empiricism. *The Philosophical Review* 60: 20-43.
- Rayamajhee, V., & P. Paniagua. [2021]. The Ostroms and the contestable nature of goods: beyond taxonomies and toward institutional polycentricity. *Journal of Institutional Economics*, 17(1).
- Samuelson, P.A. [1954]. The pure theory of public expenditure. *Review of Economics and Statistics* 36, 387-389.
- Tabarrok, A. [1998]. The private provision of public goods via dominant assurance contracts. *Public Choice* 96: 345-362.
- Vogler, J. [n.d.]. Managing social and economic externalities: industrialization and the rise of the bureaucratic state (working paper). Available at <https://janvogler.weebly.com/research.html> (accessed Feb. 22, 2024).

