Public Policy and the Private Provision of Public Goods under Heterogeneous Preferences

Mark Jacobsen, Jacob LaRiviere, Michael Price

Abstract: We compare the relative efficiency of second-best policies designed to promote the private provision of public goods. We use the provision of local public goods as our central example and discuss settings in which the model extends to choices over energy-consuming durables. We introduce preference heterogeneity by allowing a subset of agents to value the public good more than others, reflecting a form of prosocial preference. We further assume that agents face convex costs of provision, an assumption that accords well with individually provided public goods such as neighborhood amenities. We show that minimum standards are often more efficient than uniform price-based incentives in this setting. Extending our model to allow for both benefit and cost heterogeneity, we show how policy choice depends on the strength and correlation between the two forms of heterogeneity.

JEL Codes: D03, D04, H41, Q58

Keywords: Policy instruments, Public good provision

WE CONSIDER the choice among policies to increase the private provision of a public good, comparing price instruments (e.g., a subsidy to public goods) to standards (a minimum level of provision). We develop a model in the spirit of Bergstrom, Blume, and Varian (1986) and use it to explore this choice when there is heterogeneity among individuals in the economy. Specifically, we consider an economy with two types of consumers who differ in the value they place on the underlying public good and who face increasing marginal costs of provision.

While the impact of preference heterogeneity on attendant public good provision is relatively well understood, its impact on the choice of regulatory instrument and the design of first- and second-best policies remains underresearched. This is the gap in

Mark Jacobsen is at the Department of Economics, University of California, San Diego (m3jacobsen@ucsd.edu). Jacob LaRiviere (corresponding author) is at the Office of the Chief Economist, Microsoft, and the Department of Economics and Baker Center for Public Policy, University of Tennessee (jlariv@microsoft.com). Michael Price is at the Department of Economics, Georgia State University (mprice25@gsu.edu).

Received April 24, 2014; Accepted April 19, 2016; Published online February 3, 2017.

JAERE, volume 4, number 1. 2333-5955/2017/0401-0007$10.00 http://dx.doi.org/10.1086/689834

the literature we aim to fill. We motivate our model using examples from local public good provision: as in Bergstrom et al. (1986), individual agents can affect the level of local public goods such as water quality, neighborhood amenities such as green spaces, and contributions to local charities.[1] Heterogeneity in preferences means that, without regulation directed at increasing provision, the level of provision will vary across agents. We consider what this heterogeneity implies for optimal policy choice in a setting in which regulators wish to increase provision above the baseline level.

We demonstrate in our basic framework that agents with a higher value for the public good invest more in it than others. When combined with increasing marginal costs of provision, this private equilibrium involves agents with a higher value for the public good paying higher costs at the margin. As a result, the first-best policy for promoting increased provision is a Lindahl pricing scheme that induces more provision among agents with lower marginal cost; in a simplified framework all agents would provide the same level of the public good at the optimum in order to minimize total cost of provision. While theoretically appealing, such policy is operationally and informationally costly. We therefore focus in what follows on the relative efficacy of two second-best policies—uniform subsidies and minimum standards—that are easy to implement and are commonly observed in practice.[2] In our model, a uniform public good subsidy directly subsidizes provision of the public good by all agents (e.g., all agents get lawn care subsidies or rebates reducing the cost of sidewalk repair) while minimum standards require all agents to provide some minimum amount of the public good (e.g., minimum lawn care practices set by a homeowner association or minimum standards on sidewalk quality).

We find that the costs of achieving any given provision level are lower under standards than a uniform tax/subsidy, assuming all agents face the same increasing marginal cost curve for provision. The intuition for this result is straightforward: uniform taxes/subsidies induce all agents to provide more of the public good and hence preserve the wedge in the marginal cost of provision across types. Standards, in contrast, have asymmetric effects across types and serve to equalize (or, as we show, come closer to equalizing) the marginal cost of provision in the second-best. In contrast to the literature we demonstrate the possible superiority of standards even without behavioral

1. The model could also fit a variety of other public good settings such as the decision to volunteer time or provide pro bono services.

2. Homeowner associations with landscaping requirements are a common example of standards. Subsidies to sidewalk repair are paid in some cities (e.g., San Diego) while other cities (e.g., San Francisco) mandate only minimum quality levels. Martinot and Borg (1998) highlight a set of 10 similar policies in the energy sector, including subsidies to energy-efficient windows under home energy audit programs. Jacobsen, Kotchen, and Vandenbergh (2012) discuss programs that instead place minimum standards on building efficiency.

anomalies such as limited attention (Allcott, Mullainathan, and Taubinksy 2012) or temptation (Tsvetanov and Segerson 2014).[3] Our results highlight that for any fixed level of provision, the welfare advantage of standards is (i) increasing in the weight the high-valuation types place on the returns to providing the public good and (ii) single-peaked in the proportion of such types in the economy. Importantly, though, this result relies on identical increasing marginal costs faced by all agents.

Next, we extend the model to allow heterogeneity in the cost of provision in addition to heterogeneity in preferences for the public good. We show that the relative superiority of standards is typically eroded as cost heterogeneity becomes large, with policy choice depending on both the degree and correlation of the two types of heterogeneity. We thus urge caution in interpreting our results: we find that standards will be more efficient than uniform prices (in lieu of agent-specific first-best Pigouvian taxes) in certain cases but more generally that the two types of heterogeneity will compete to determine the best policy.

Our consideration of preference heterogeneity in the context of public goods accords well with prior studies showing that individuals differ in their willingness to pay for public goods. Heterogeneity in provision of local public goods, such as landscaping around private homes, is self-evident in any suburban neighborhood.[4] Laboratory and field experiments also demonstrate significant variation in public good provision (Dana, Cain, and Dawes 2006; Dana, Weber, and Kuang 2007; Lazear, Malmendier, and Weber 2012; DellaVigna et al. 2013; Andreoni, Rao, and Trachtman 2017).

While we develop the theory in a public goods setting, under additional assumptions our work can also complement the large body of research on first- and second-best regulatory instruments for reducing environmental externalities (e.g., Weitzman 1974; Bovenberg et al. 2008; Fowlie and Muller 2013; Jacobsen 2013; Carson and LaRiviere 2017). Because our model considers provision of a single public good, we limit our discussion here to cases in which the only elastic margin for reducing energy use is in the choice of an energy-using durable.[5] This abstracts from settings in which a second, utilization, choice influences the externality and could interact with both policy and preferences, most typically favoring price-based instruments that can act on both margins

3. In the context of second-best policies designed to reduce deadweight loss from externalities, standards may also be preferred to taxes in the presence of compensation requirements (Bovenberg, Goulder, and Jacobsen 2008).

4. Tiebout sorting can manifest directly as a result of preference heterogeneity over local public goods.

5. Households face choices over energy efficiency levels that embed an amount of carbon generated, therefore providing environmental quality as a public good. In the United States, household-level choices over energy use account for approximately 40% of all greenhouse gas emissions.

simultaneously. We discuss refrigeration and lighting below as examples in which the product choice margin is likely most important in determining the externality, though we note that some intensive use elasticity may still be present.[6] We show how the optimal policy depends critically on the motivation for prosocial behavior: If "green" agents choose energy-saving durables because of lower cost (e.g., because the agent receives utility from being an early adopter or a warm glow from personally providing the good), the resulting cost heterogeneity means that price-based policy will dominate standards, a familiar result in this literature. However, we also show that in certain cases, preference heterogeneity—and therefore a policy involving minimum standards—could dominate instead. This requires that preferences for the public good differ sharply enough that individuals provide different amounts even though their contributions produce only tiny changes in the aggregate. We show how the social efficiency preferences described in Charness and Rabin (2002) produce this effect; these preferences scale up an individual's value for the public good along with population. The opposing policy implications depending on the underlying motivation of agents suggest the importance of further empirical work before applying the model to settings with large numbers of agents.[7]

Our approach differs from the existing literature on public goods along a number of important dimensions. First, we take the perspective of optimal second-best policy, comparing two broad classes of second-best policy (price incentives or mandates applied uniformly) aimed at increasing the private provision of public goods. The literature discussing choice of prices or mandates has more typically built on the externalities model (as in, e.g., Weitzman 1974; Segerson 1988; Hoel and Karp 2001; Bovenberg et al. 2008; Holland 2012) and treated preferences over the public good as uniform with heterogeneity instead affecting the cost of provision. Our public goods model is a special case of an externality problem in which each agent's utility is a function of the sum of all agents' contributions to the public good.[8] While we cannot make general statements on the externalities model, this does allow us some extensions and potential

6. Prior work explores behavior along multiple margins—e.g., the choice of fuel economy and subsequent miles driven—and how optimization errors interact to influence policy mix (e.g., Allcott et al. 2012). Our results apply even in the absence of such interactions, but this does limit direct application of our model to cases in which adjustments along the intensive margin are small.

7. A variety of studies demonstrate substantial heterogeneity in provision of energy conservation (e.g., Kotchen and Moore 2007; Saphores et al. 2007; Jacobsen 2013) or willingness to take costly actions to conserve resources to improve environmental quality (e.g., Allcott 2011; Costa and Kahn 2013; Ferraro and Price 2013; Ito, Ida, and Tanaka 2013; Metcalfe and Dolan 2013), but the form of the motivation in utility is not studied specifically.

8. We note, though, that this particular additive structure is used in much of the externality literature.

applications in the space of policies directed at problems such as externalities from aggregate energy use. Our theoretical approach remains grounded in the public goods setting and builds most closely on the Bergstrom et al. (1986) model of public good provision in which marginal utility of provision of the public good decreases at different rates across heterogeneous agents.

Second, our model allows for a more general preference structure than prior work exploring the private provision of public goods. The key feature of our model is that agents have heterogeneous values for aggregate provision of the public good. This can appear through a variety of mechanisms explored in the literature including agents who (i) exhibit heterogeneous concerns for efficiency or overall public good provision (Becker 1974; Charness and Rabin 2002), (ii) face social pressures when forming preferences for the level of overall public good provision (Levine 1998; Akerlof and Kranton 2000; DellaVinga, List, and Malmendier 2012), or (iii) make optimization errors or succumb to temptation when valuing overall public good provision (Allcott et al. 2012; Allcott and Taubinsky 2014; Tsvetanov and Segerson 2014). We also consider the implications of a different type of preference in which individuals receive utility benefit associated only with their own contribution (irrespective of the overall level) that appears as warm glow, social identity, or prestige (Andreoni 1989; Hollander 1990; Glazer and Konrad 1996; Kessler and Milkman 2014). We show how this type of heterogeneity manifests as cost heterogeneity for the agent and competes with heterogeneity in preferences for the overall level of the public good in determining which policy will be optimal.

The remainder of the paper is organized as follows: Section 1 introduces the theoretical model and derives analytical results. Section 2 provides a simple algebraic example of the model, highlighting the trade-off between the two types of heterogeneity in determining policy. Section 3 discusses application of the theory to local public goods and extensions to the context of refrigeration and lightbulb efficiency. Section 4 presents conclusions.

1. PUBLIC GOODS PROVISION AND INDIVIDUAL PREFERENCES

We begin by outlining a theoretical model of public goods provision in a world in which agents face identical and linear costs of provision but have heterogeneous preferences for the public good. This is the baseline model analyzed in Bergstrom et al. (1986) and provides us a framework to analyze the cost of alternative policies designed to promote increased provision. We then extend the model such that agents face increasing marginal costs of provision. Increasing marginal cost is relevant to a broad range of public goods provided by individual households directly and is the key focus of our theoretical results. We begin with the linear costs model to make salient the differences (from the perspective of determining optimal policy) that emerge when we introduce increasing marginal costs. Section 1.4 goes on to extend the model to allow a second form of heterogeneity: we relax the assumption of identical costs and allow heterogeneity both in costs and in preferences. We explore how the two sources of heterogeneity compete in determining optimal policy choice. Proofs appear in appendix A.

1.1. Modeling Preference Heterogeneity

We define utility over a numeraire private good, c, and a public good, X. Provision of X will be the sum across N members of the economy, where x_i is the contribution of the ith member, $\sum_{i=1}^{N} x_i = X$. First, following Bergstrom et al. (1986), individuals provide the public good subject to a constant marginal cost, p, such that the total cost of providing x_i units of the good is px_i.

We consider the provision of a public good such as donations to or time spent working at a local charity, purchases of landscaping services or improved sidewalks, or in later extensions the purchase of a more expensive product that reduces pollution. In these latter examples we model the choice of a given product as embedding pollution reduction and thus providing a public good: environmental quality. For example, the decision between two refrigerators that are identical in every way except that one uses less electricity than another presents one way to provide public goods associated with lower electricity use. In this regard, the extension of our model to energy use externalities is limited to decisions along a single margin of durable choice without regard to later utilization decisions.

Preferences over the two goods will be given by

$$\begin{aligned} U_i(c_i, X) &= c_i + \widetilde{V}_i f(X) \quad \text{s.t. } y_i = c_i + px_i \\ \Rightarrow U_i(c_i, X) &= y_i - px_i + \widetilde{V}_i f(X). \end{aligned} \tag{1}$$

In equation (1) utility is assumed linear in the numeraire good and weakly concave in the agent's valuation of the public good. As is standard in the public goods literature, we will assume that $\lim_{x \to 0} f'(X) = \infty$ and $f'(X) \geq 0$ for all X. In the model, x_i is actual units of the public good. For example, in the context of landscaping, it could measure how well manicured a house's lawn is. In the context of refrigerators it would be the increase in net present cost of owning a particular refrigerator (purchase price plus energy use) that is associated with lower pollution. For now we assume the net costs are the same no matter which household provides the cleaner refrigerator; households that under- or overvalue future energy costs in different ways will appear in our examples of cost-side heterogeneity in section 1.4 below.

Consistent with Bergstrom et al. (1986), this modeling approach implies that the marginal benefit agents receive from contributing to the public good will be a function of extant levels of X provided by other agents in the economy. We show in appendix B that this formulation is equivalent to a model with decreasing marginal utility over the consumption good and a linear budget constraint, such that our results are not driven by the quasi-linear functional form assumption. The appendix follows existing work on the private provision of public goods (e.g., Bergstrom et al. 1986; Andreoni 1989).

The heterogeneity we consider in our model arises through the term $\widetilde{V}_i$, which scales each agent's preference for the public good.[9] For simplicity we will allow two types of agents, a share α with $\widetilde{V}_i > 1$ and a share $(1 - \alpha)$ with $\widetilde{V}_i = 1$.

There are several interpretations for heterogeneity in $\widetilde{V}$. First, the differences could be due to strict neoclassical preference heterogeneity. Indeed, there is significant field evidence that agents have varying preferences for privately provided public goods (Kotchen and Moore 2007; Saphores et al. 2007) or attitudes toward different strategies they may take to provide more environmental quality (Poortinga et al. 2003).[10] Second, variation in $\widetilde{V}$ could embed heterogeneity in marginal benefits for the public good from any source as long as it enters multiplicatively (e.g., social norms as in models described by Benabou and Tirole [2006] or DellaVigna et al. [2012]).[11]

Third, heterogeneity in $\widetilde{V}$ nests a number of preference structures explored in the prior literature. The most important for our paper here are efficiency preferences modeled in Charness and Rabin (2002). The efficiency preference interpretation in Charness and Rabin's study provides a strong link between preferences here and the earlier behavioral literature. They consider the following utility structure:

$$U_i(c_i, X|\Theta) = c_i + f_i(X) + e\sum_{j \neq i} f_j(X).$$

Added to the standard neoclassical specification is a term that includes the sum of all other agents' valuation of the public good (multiplied by a weight e). This implies that individual i will adjust their own provision of the public good toward social efficiency, while an agent without this added term will not. In our setting this is equivalent to set-

9. We assume that the shape of preferences for the public good given in $f(X)$ is common across agents. This assumption provides transparency in the analysis but is not necessary in that heterogeneity in f would also produce our result to the extent it introduces differences in marginal benefit across types.

10. There is also an immense body of evidence from laboratory experiments suggesting heterogeneity in the willingness of individuals to contribute to a group fund. For example, Brandts and Schram (2001) explore behavior in a setting in which subjects report a contribution function that states contribution levels for various rates of transformation between public and private accounts (returns to the public good). They find that some subjects behave in accordance with a utility function defined solely over own earnings while others behave in a manner consonant with social preference or other-regarding behavior. This closely matches our setting here. Fischbacher and Gaechter (2008) identify substantial heterogeneity in preferences and discuss the decline in public good provision over time. Kurzban and Houser (2001) use a circular public goods game and classify individuals into three distinct types whose underlying motives for giving differ.

11. Similarly, Krupka and Weber (2013) elicit beliefs about the social appropriateness of various allocation decisions in variants of the dictator game and show similar heterogeneity in perceived norms for giving.

ting $\widetilde{V} = 1 + e(N - 1)$. We allow heterogeneity in that some agents have the prosocial term while others do not. Note that this differs from models in which utility benefits accrue from the act of giving in and of itself and so would be a function of x_i rather than X.[12] Appendix C shows how this preference structure works to maintain differences in preference for the public good even as N grows large. We focus first on local public goods in which N is relatively small and this particular form is not needed.

Finally, heterogeneity in $\widetilde{V}$ could reflect optimization errors that arise through limited attention or temptation as in Allcott et al. (2012) or Tsvetanov and Segerson (2014). We focus on the differences created in the final value an agent places on the public good but note that the form of the model generating that difference could imply very different welfare analyses and different aggregate levels of provision at the social optimum. We focus only on the relative efficiency of the two second-best policies above, holding fixed the level of public good provided.

Assuming an interior solution, private agent optimization leads to first-order conditions characterizing the following Nash equilibrium:

$$p \geq \widetilde{V}_i f'(x_i^* + X_{i\neq j}^*) \quad \text{with equality if } x_i^* > 0. \tag{2}$$

There are two important features of this expression. First, the marginal cost of providing an additional unit of the public good, p, is identical for every agent in the economy. Second, in the private equilibrium, it is quite possible that only the high-value agents provide positive levels of the public good. Specifically, as $\widetilde{V}_i$ increases, $f'(x_i^* + X_{i\neq j}^*)$ must decrease so that their product equals the price of provision p. Heterogeneity in $\widetilde{V}_i$ gives rise to the free-rider problem: agents with a larger value for the public good provide it while the others free ride. As a result, the agents with lower valuation might not provide any of the public good.

From a policy perspective there are two parts to the first-best solution. First, a regulator must find the optimal level of the public good, X^*. The optimal level will set the marginal value of the public good across all agents in the economy equal to marginal cost such that $p = \sum_i \widetilde{V}_i f'(X^*)$. Second, the regulator must distribute the cost of provision across agents in the economy. Most generally this can involve a Lindahl price scheme in which each agent pays for public good provision proportional to their own marginal benefit. In this simple model of quasi-linear utility and constant marginal cost, if the social welfare function sums agents' utility with equal weight, the allocation of provision does not matter for social welfare. Put another way, once the social planner decides X^*, a symmetrically weighted total social welfare function takes the same value

12. We also do not directly nest relativistic preferences in which one agent's level of provision relates to other agents' level of provision, though we believe our results will extend and discuss this situation below.

regardless of how aggregate provision is allocated across agents. This result is driven by the identical, and linear, costs of provision. As a result, asymmetric Lindahl pricing is not needed to achieve a first-best outcome, and a uniform subsidy to provision sufficient to reach X^* will be first-best. Similarly, a minimum standard inducing all agents to provide X^*/N units of the public good would also achieve a first-best outcome.

1.2. Increasing Marginal Costs

We now extend the model such that, in place of linear costs, individuals provide the public good subject to a strictly convex cost function $h(x_i)$ with $h'(x_i) > 0$ and $h''(x_i) > 0$. This assumption accords well with household-level provision of local public goods such as landscaping, where increases in effort produce decreasing amounts of the public good. Our later examples of energy conservation and the associated private provision of environmental quality also fit an increasing marginal costs framework: For example, in the context of energy conservation provided through lighting choice, the switch from incandescent to compact fluorescent (CFL) bulbs may be relatively cheap. However, if an individual wishes to provide even more conservation, they may switch to more expensive light-emitting diode (LED)-based products, effectively moving out a rising marginal cost curve. Similarly, electricity-intensive durable goods such as refrigerators increase in cost at an increasing rate as energy efficiency of the product rises.[13]

Preferences over the two goods, now including convex costs $h(x_i)$, will be given by

$$\begin{aligned} U_i(c_i, X) &= c_i + \widetilde{V}_i f(X) \quad \text{s.t. } y_i = c_i + h(x_i) \\ \Rightarrow U_i(c_i, X) &= y_i - h(x_i) + \widetilde{V}_i f(X). \end{aligned} \tag{3}$$

The standard first-order conditions for the privately optimal provision of the public good are given in our model by

$$h'(x_i^*) \geq \widetilde{V} f'(x_i^* + X_{i \neq j}^*) \qquad \forall \ i, \tag{4}$$

with equality if $x_i > 0$ and where $X_{i \neq j}^*$ is the sum of all agents' $j \neq i$ privately supplied optimal levels of the public good. Unlike the linear cost of provision case considered above, with convex costs the marginal cost of provision will differ across agents depending on how much of the good they elect to provide. As a result, differences in the marginal benefits from the public good (i.e., heterogeneity in $\widetilde{V}$) across agents create differences in the costs of provision across agents at the margin (e.g., evaluated at x_i^*).

13. Our assumptions in this simple version of the model will be met to the extent consumers face similar costs in the marketplace and the technologies are close substitutes in providing a final good (e.g., refrigeration to a certain temperature or lumens of light).

A critical component of the model with increasing marginal costs is based on the level of provision of the two types of agents in the private equilibrium. Lemma 1 summarizes these differences:

Lemma 1: Agents with higher valuation of the public good provide more than those with lower valuation.

The intuition is straightforward: since the marginal benefit associated with each unit of the public good is higher for some agents, they will always provide more of the public good than those with lower valuation. Moreover, given our assumption of a common underlying cost structure, this creates a wedge in the marginal cost of provision. Figure 1 shows this graphically and illustrates the privately optimal levels of provision, $\hat{x}_g$ and $\hat{x}_u$, for agents with higher and lower valuation, respectively.

Next, we consider the cost-minimizing allocation rules. Assuming increasing marginal costs, the model immediately leads to the following:

Proposition 1: For any level of public goods provision $\widetilde{X}$, it is cost minimizing to have all agents provide identical quantities.

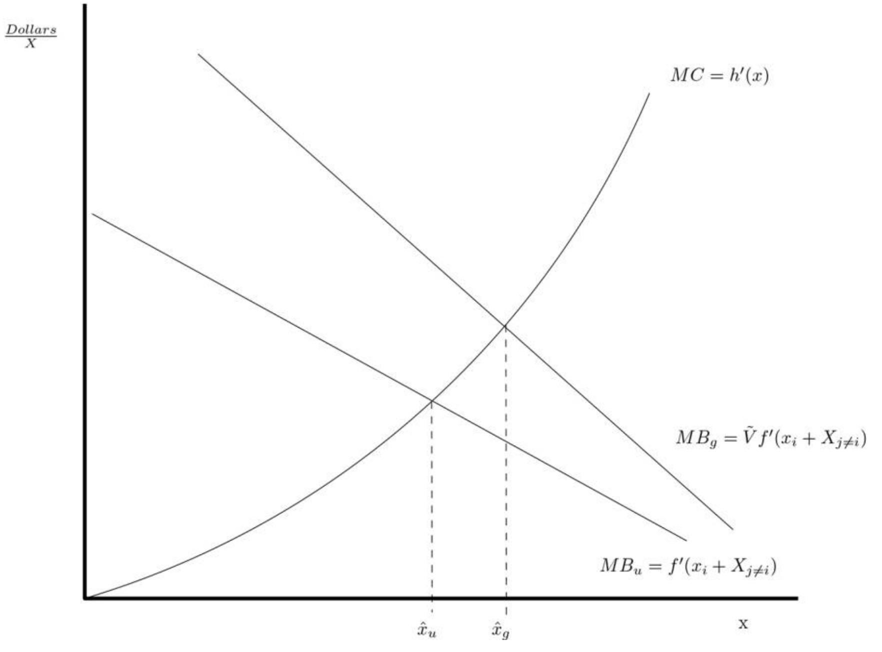

Figure 1. Basic equilibrium

Corollary 1: Private provision of the public good is socially efficient if all agents value the public good as would the social planner.

Proposition 1 is simply the classic equimarginal principle in the context of our model. Since the marginal cost of provision is increasing and symmetric across agents and the benefits to increasing the aggregate $\widetilde{X}$ do not depend on the identity of the contributor, the minimum-cost provision of $\widetilde{X}$ will involve $x_i = x_j$. However, as noted in figure 1, the marginal cost of provision for some agents is larger than for others. As a result, the private equilibrium cannot minimize cost since there is a wedge in marginal costs across agents. Hence, reallocating production of the last unit produced by a high-cost agent to a lower-cost agent would lower the overall costs of provision.

Proposition 1 implies that the private equilibrium in this model is inefficient since the two types of agents provide different levels of the public good and thus have different costs of provision on the margin. There are two assumptions driving this result. First, we must assume preferences for the public good are heterogeneous. However, as shown in the previous section, that assumption in and of itself is not sufficient to create violations of the equimarginal principle if the cost of provision is linear. The differences in the marginal cost of provision across types require combination with a second assumption—increasing marginal costs of provision. Those two assumptions, both of which accord with empirical evidence, produce heterogeneous costs of provision on the margin between the two types of agents.

Corollary 1 includes two important features that make it different from the classic definition of equilibrium public goods provision (e.g., in Bergstrom et al. 1986). First, each agent provides their own contribution of the public good through independent convex cost functions, $h(\cdot)$. These agents must therefore have the same preferences in order to guarantee provision of identical quantities as required for cost minimization. Second, these preferences must place enough weight on the public good to bring the aggregate level of provision up to the efficient level. Note that the social planner here implicitly considers only direct benefits of provision, contained in $f_i(X)$.[14] We will abstract from the optimal level of public good provision in what follows by considering the planner's problem as a cost minimization subject to reaching a given level of provision in aggregate.

1.3. Preference Heterogeneity and Policy Choice

We now turn to the motivating question about cost-minimizing policy. In this subsection and the remainder of the paper we maintain the assumption of increasing marginal

14. If the planner also considers feedbacks in the efficiency preference term (i.e., accounting for $s > 0$ as opposed to only $f_i(X)$), it leads to an increase in the optimal level of public good provision but will not affect our results below, which focus on cost minimization.

costs of provision. We first describe the optimal policy: Lindahl pricing. Given the practical difficulties in implementing individual-specific policies for public good provision, we focus most on second-best policies that are not differentiated across individuals.

The first-best Lindahl prices set the marginal price of the public good for each agent such that the absolute value of the price ratio is equal to the social marginal rate of substitution between the public good and the private good. This is the pricing scheme developed in Bergstrom et al. (1986). The increasing marginal costs in our model mean that the prices (subsidies in this case) must simultaneously account for differences in valuation of the public good and rising costs of provision. Specifically, the first-best Lindahl pricing scheme will set agent-specific subsidies, s_i, such that $\widetilde{V}_i f(X^*) + s_i = h'(X^*/N)$, where X^* is the first-best level of the public good. Importantly, increases in $\widetilde{V}_i$ are associated with decreases in s_i since marginal cost $h'(X^*/N)$ will be constant across agents at the efficient point.

Proposition 2: There is an asymmetric price instrument that leads to efficient provision of any level of the public good, $\widetilde{X}$. Under such a policy, the subsidy for agents with low valuation, τ_u, is strictly larger than the subsidy for agents with high valuation, τ_g.

The intuition for proposition 2 is related to the classic Lindahl price solution in public goods problems. The difference here comes from our assumption of convex costs, which necessitates equal provision of the goods across types for efficiency. An asymmetric price policy can achieve this. Under any such policy, the social planner uses targeted subsidies to shift each agent's marginal benefit curve such that it intersects the marginal cost curve at the same point.[15] The resulting equilibrium would entail each agent providing the same level of the public good at the same cost, thereby making the provision efficient.

However, this policy would be quite difficult to implement in most settings because of the difficulty in accurately identifying each type.[16] Those with high valuation for the public good would have an incentive to imitate low-valuation types, removing the corrective effects of the Lindahl pricing and returning the model to a case in which all agents face the same cost curve for provision.

We therefore compare the performance of two second-best policies, both of which are observed commonly in practice: (i) uniform price instruments (subsidy to public

15. This is similar to the setting in Diamond (1973), where different agents produce different externalities and so would require targeted Pigouvian taxes or subsidies.

16. Resale trade could also lead to inefficient levels of public good provision. Intuitively, the problems would be akin to those faced by a third-degree price discriminator under resale, though we leave questions of strategic resale and third-degree price discrimination to future work.

good provision) and (ii) standards that enforce a minimum level of public good provision.[17] For simplicity, we assume perfect enforcement of both policies and we consider policies designed to get the economy to a fixed level of the public good. All analysis preserves the increasing marginal cost of provision and heterogeneous benefits assumptions. In the next subsection we extend the model to also allow the costs of provision to vary across agents.

Understanding the mechanics of how uniform subsidies or standards in our model map to practical policies is important: subsidies to landscaping, sidewalks, or efficient durables (e.g., through rebate programs) appear directly as a price policy in our model; any agent may receive the rebate if they choose to supply more of the good. Minimum standards, in contrast, change the choice set, making actions that do not provide a basic level of the public good illegal or impossible, for example, a minimum frequency of lawn maintenance or minimum allowable efficiency ratings for durable goods. These policies do not reward agents for going above the minimum standard (the key difference for the purposes of our model), but they do make it difficult or impossible to choose a level of provision below the standard.

The relative efficiency of a minimum standard versus a uniform price instrument is summarized in the following two propositions:

Proposition 3: For any level of regulated public goods provision $\widetilde{X}$ such that the standard binds for all agents, $\widetilde{X}/N \geq \hat{x}_g$, a standard is always more efficient than a uniform price instrument.

Proposition 4: For any level of regulated public goods provision $\widetilde{X}$ such that the standard binds for agents with low valuation but not for those with high valuation, $\hat{x}_g \geq \widetilde{X}/N \geq \hat{x}_u$, a standard is always more efficient than a uniform price instrument.

The intuition behind proposition 3 follows from our assumption of a common cost function. If the standard binds for all agents, then they all provide $\widetilde{X}/N \equiv \bar{x}$ units of the public good and reach aggregate provision at minimum cost. A uniform price subsidy, on the other hand, preserves a wedge in the level of provision between types. The difference in provision, and therefore marginal costs across types, means that the price

17. Note that in our setting standards placed at the producer level could still be passed through to consumers (the agents in our model) as price effects. For example, an emissions quota applying to electricity production will appear as a higher electricity price for all households; from the perspective of the agents in our model, it appears as a price instrument. An emissions tax would similarly manifest as a price instrument to consumers. Conversely, a technology standard limiting the types of goods that may be produced would eliminate some products from our agents' choice set, constraining them to a minimum level of provision.

policy will always be less efficient than a standard as long as all agents face identical increasing marginal cost functions. We show below that when we relax the identical costs assumption, the advantage of standards can be eroded or overturned depending on the degree of cost heterogeneity. Finally, as discussed above, subsidies targeted by type can still produce efficiency, though we will argue that these are infeasible in most situations and share an important similarity to standards in our context; they have a greater impact on the decision of those with low valuation for the public good.

Although the mechanics are more complex, proposition 4 results from the same mechanism: subsidies preserve a greater wedge in the marginal cost of provision across types. Since the standard in this case is not binding on all agents, some will provide more of the public good than others under either policy. The key intuition for the result is that the subsidy increases provision from both types of agents while the standard increases provision only from agents with low valuation who also have lower marginal costs of provision. The concavity of benefits from the public good introduces an indirect effect reinforcing this result: if, because of declining marginal benefits, the unconstrained agents provide somewhat less of the good, the wedge between the two types of agents will only be further reduced. The standard is therefore unambiguously preferred for any level of provision greater than the privately optimal level.

The effect of the price instruments and standards considered by proposition 4 is shown graphically in figures 2 and 3. In figure 2, the subsidy for public good provision shifts the marginal benefit curve for both types up by the level of the subsidy, τ. Each type sets private marginal benefit (now including the subsidy) equal to marginal cost, preserving a wedge in the marginal cost of provision across types: $MC(x_u(\tau)) < MC(x_g(\tau))$. The size of the wedge is directly related to the inefficiency of the price instrument: on the margin it would be cheaper to have those with low valuation provide more of the good. As a result a price instrument cannot be efficient because it preserves this wedge. Figure 2 abstracts from the Nash equilibrium with best-response functions in order to highlight this first-order effect. A more precise graph would also show the indirect effects (which we consider fully in the propositions) coming from extant provision by other agents. For reference we do include indirect effects in figure 3.

Figure 3 shows the effect of a standard that binds only for agents with low valuation. In this case, they must provide at a level greater than their private optimum ($\bar{x} > \hat{x}_u$). This brings the marginal cost of provision for the two types closer together, holding the choice of the high-valuation agent fixed. As a result, the wedge in marginal costs across types is lower with a standard than with a price instrument. Since the wedge in marginal cost across types is lower, it is less costly to provide the same level of the public good.

In both figure 2 and figure 3 there are possible indirect effects of the policy instrument that also act to increase the relative efficiency of the standard. In both cases aggregate provision of the public good is rising, which tends to lower the marginal benefit of additional private provision. In the case of the standard, provision of agents with low

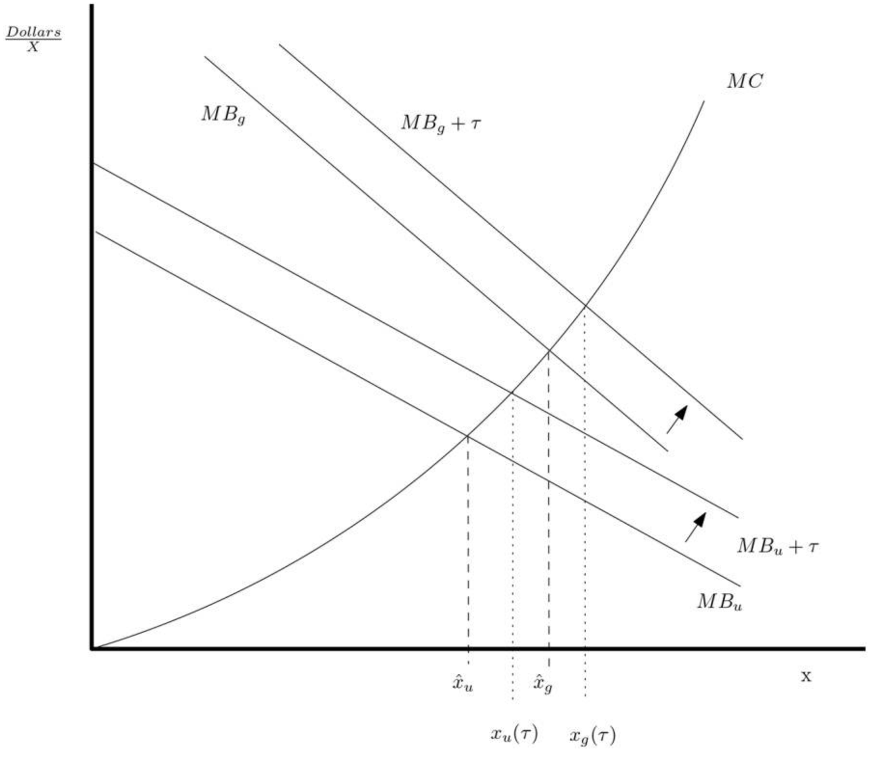

Figure 2. Equilibrium with a price instrument

valuation is fixed (at the level of the standard), so the indirect effect means the high-valuation agents will provide (weakly) less than before ($x_g(\widetilde{X}) \leq \hat{x}_g$). The direct and indirect effects of the standard then both act to narrow the wedge in marginal costs between agents. Under a price instrument, in contrast, indirect effects will act on both types simultaneously and therefore preserve the costly wedge in provision.[18]

In addition to demonstrating the advantage of a standard, the model also permits consideration of the relative size of this advantage with respect to two key parameters: the proportion of prosocial agents and the relative strength of their preferences. Corollaries 2 and 3 below are the first steps toward the comparison and show how provi-

18. In both cases, the indirect effect arises from the public good modeling assumption: the amount of the good provided by other agents affects the marginal benefit from provision by any given agent. If agents ignored the marginal benefit from their own contribution to the public good, the indirect effect above would not occur. We note that for both standards and taxes the indirect effects are second-order to the direct effect coming from the wedge in marginal costs.

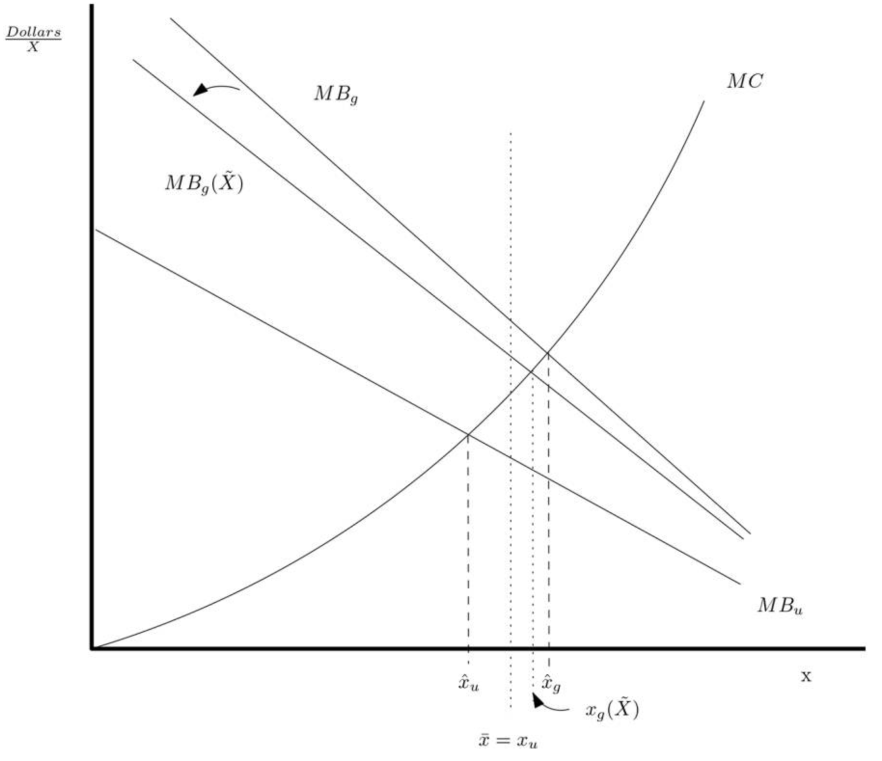

Figure 3. Equilibrium under a standard

sion of each type of agent changes as a function of α, the proportion of prosocial agents, and $\widetilde{V}$, the strength of their preferences under the price instrument:

Corollary 2: The provision of prosocial agents in the case of price instruments, x_g^t, and the price instrument/subsidy needed to attain a given level of public good provision, τ, are both always decreasing in the percentage of prosocial agents, α.

Corollary 3: The provision of prosocial agents in the case of price instruments, x_g^t, is increasing, and the price instrument/subsidy needed to attain a given level of public good provision, τ, is decreasing in the strength of the prosocial preference, $\widetilde{V}$.

The intuition behind corollary 2 is straightforward. The first part of the corollary follows directly from our assumption on concavity of $f(X)$; displacing one low-valuation agent with a high-valuation counterpart leads to an increase in aggregate production. At prevailing provision levels, we thus have that the marginal cost exceeds marginal benefit. In equilibrium, this implies that each agent will provide less of the public good.

The intuition for the second part of the corollary is as follows. The unregulated level of public good provision is directly related to the proportion of prosocial agents in the economy. As the proportion of such agents increases, the change in provision required to achieve any target level of provision falls. By convexity of the marginal cost function, the resulting subsidy needed to achieve the desired change in provision falls. The intuition behind corollary 3 is almost identical, except that now the increase in unregulated provision comes through increased strength of individual preferences in $\widetilde{V}$ rather than an increase in the number of prosocial agents overall.

The corollaries lead to two results on the size of the cost advantage offered by a standard:

Proposition 5: For any level of regulated public goods provision, $\widetilde{X}$, such that the standard binds on all agents, $\widetilde{X}/N \geq \hat{x}_g$, the difference in welfare between the two policies, Δ_{ts}, is single peaked in the percentage of prosocial agents, α.

The intuition underlying proposition 5 is again straightforward. Consider first the case in which all agents are of a given type. In this case, the two policies are equivalent and there is no difference in welfare. However, the welfare effects of the two policies begin to diverge once we allow for heterogeneity in preferences. In such a world, agents provide different levels of the public good at different marginal costs under the price instrument but face identical costs under a binding standard.

Corollary 4: For any level of regulated public goods provision $\widetilde{X}$ such that the standard binds for all agents, $\widetilde{X}/N \geq \hat{x}_g$, the difference in welfare between the two policies, Δ_{ts}, is everywhere increasing in the strength of the prosocial preference, $\widetilde{V}$.

The intuition for corollary 4 is as follows: By convexity of the cost function, the welfare gain that arises when reallocating one unit of provision away from the prosocial agent is greater the greater the initial wedge in costs. As shown above, stronger prosocial preference, $\widetilde{V}$, induces more provision from prosocial agents and increases the wedge in costs between the two types. As a result, the tax leads to progressively greater costs than the standard as the strength of the prosocial preference increases.

Taken together, the results in this subsection consider the effect on public goods provision when a fraction of people in the economy have prosocial preferences. We show how such types provide more of the public good in an unregulated setting, leading to a wedge in the marginal cost of provision between high-valuation agents and others in the economy. This wedge is preserved if the government employs a uniform price instrument to promote increased provision of the public good. Standards, in contrast, have a greater impact on low-valuation agents and thus serve to reduce the cost wedge and lower the costs of obtaining any given level of provision. Finally, we find that the relative benefit of standards over price instruments is increasing in the degree of preference heterogeneity in the economy up to the point at which prosocial agents constitute 50%. In our model, this comes through two channels: stronger preferences among individual prosocial agents or an increase in their overall number.

1.4. Asymmetric Costs

We now consider asymmetries across types in the cost of providing the good. To varying degrees, cost asymmetries are quite likely to exist in many important policy settings; we discuss these in the specific examples below. We show how the two types of heterogeneity push in opposing directions on efficiency of the policies but do not take a stand on the relative magnitude of cost versus benefit heterogeneity in general.

To explore cost asymmetries, we assume that prosocial agents, in addition to having the $\widetilde{V}$ parameter governing the strength of their preferences, also have a parameter δ differentiating their marginal costs of abatement. The prosocial agent's maximization problem thus becomes

$$\begin{aligned} U_g(c_g, X|\Theta) &= c_g + \widetilde{V}f(X) \quad \text{s.t. } y_g = c_g + \delta h(x_g) \\ \Rightarrow U_g(c_g, X|\Theta) &= y_g - \delta h(x_g) + \widetilde{V}f(X), \end{aligned} \tag{5}$$

where values of δ greater than one scale up the cost of abatement for high-valuation agents (making cost positively correlated with prosocial preferences), and values less than one scale it down (introducing a negative correlation).

Allowing for cost heterogeneity is important for three main reasons. First, the act of providing in itself may create utility that an individual does not get from increases in aggregate provision by others. This form of individualistic warm-glow preference will lower the utility cost of providing the public good for some agents but not others. Second, heterogeneous private benefits that are associated with the public good also manifest as cost heterogeneity for the same reason. For example, if the private value to a household of having their own lawn be well manicured is large, this effectively decreases the cost of lawn care for that individual, separate from any valuation of the aggregate public good. Third, in the context of energy-using durables, individuals may value future reductions in energy use differently (e.g., because of inattention or different discount rates), making the net costs of a more efficient appliance different, and even negative if the energy savings are unexpected. Each of these is embedded in the model as a source of heterogeneity in individual costs of provision.

The extension to cost heterogeneity leads to the following two propositions:

Proposition 6: With asymmetric costs, the amount of public good provided by prosocial agents is inversely related to their relative cost position. The (uniform) subsidy needed to reach any level of public goods provision is increasing in the relative marginal cost of provision for prosocial agents.

Proposition 6 states two very intuitive results. First, as the cost of provision for prosocial agents increases, they provide less of the public good. Second, holding the cost of provision for other agents constant, the subsidy required to induce any given level of public good provision is increasing in the cost of provision for prosocial agents.

Proposition 7: With asymmetric costs, the relative efficiency of a uniform price instrument vis-à-vis a standard is greater when the costs for prosocial agents are negatively correlated with the strength of their preferences.

Proposition 7 states that the relative benefit of a standard over a price instrument falls if prosocial agents have a lower cost of provision. Conversely, the relative benefits of standards over taxes can rise if prosocial agents have a higher cost of provision and the standard binds only for agents with low valuation. The intuition for these results is shown in figure 4. The figure provides an example in which provision in the unregulated equilibrium is in fact efficient since the prosocial agent's marginal cost of abatement curve is sufficiently less than the costs of the other agents. In this example, it is easy to see how a uniform subsidy can maintain efficiency by shifting both types' provision up. A standard, in contrast, will introduce a wedge in marginal costs leading to a reduction in social welfare.

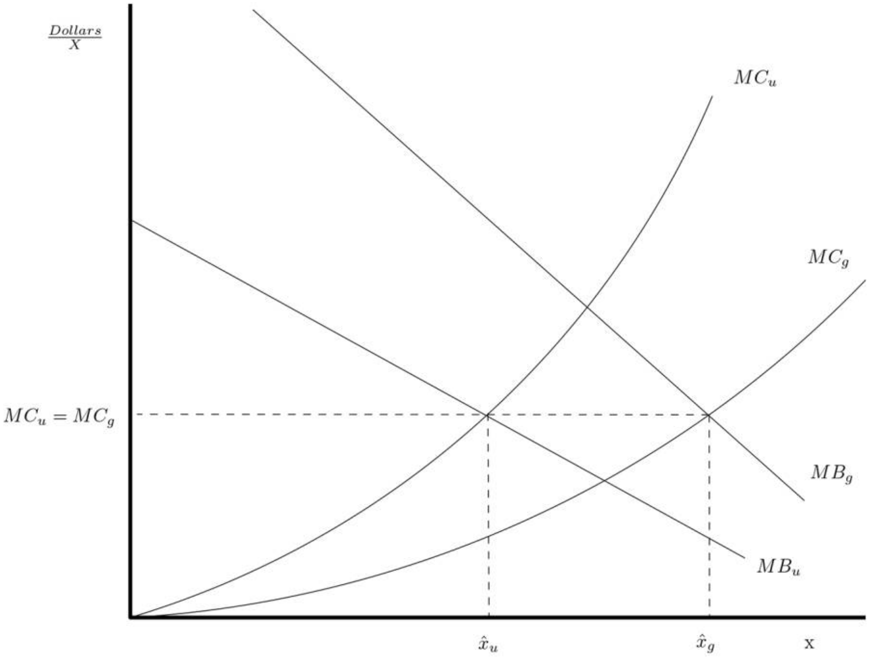

Figure 4. Case of asymmetric costs of abatement

More broadly, cost heterogeneity will tend to decrease the advantage of a standard relative to a subsidy. For a given level of benefit heterogeneity, if the magnitude of the cost heterogeneity is sufficiently large, these results imply—although do not prove—that the subsidy policy will dominate the standard. The extreme case makes the intuition clear: consider the situation in which benefit heterogeneity is infinitesimally small and cost heterogeneity is large. The model then reduces to a classical public goods problem with heterogeneous costs in which the price instrument dominates. Below, we provide a more precise depiction of these competing effects and how they affect the relative superiority of standards using a simple quadratic approximation.

2. A SIMPLE QUADRATIC EXAMPLE

While the key inequalities appear in the general model above, there is also important intuition in a parsimonious analytical example. A simple quadratic form makes the key elements in our model clear and allows us to explicitly compare the trade-off between heterogeneity in preferences and heterogeneity in costs. We use the quadratic approximation to solve explicitly for the advantage of a standard relative to a price-based instrument and investigate how each factor enters. We then add heterogeneity in costs and again compare the competing effects influencing the optimal choice of policy instrument.

In what follows, we assume that the marginal benefit from provision of the public good is constant and that the marginal cost of provision rises linearly from the origin. Although this imposes a particularly strong structure, we believe that such a parameterization fits well with a number of policy-relevant scenarios. For example, reductions in carbon emissions in a particular country and year (now treating countries as the agents in our model) are likely to fit the constant marginal benefits case closely.[19] Linearly increasing marginal cost provides a quadratic approximation to public goods produced with convex costs to individual agents.

Our example then defines

$$f_i'(X) = m \text{ (constant marginal benefits)},$$
$$h'(x_i) = bx_i \text{ (marginal cost rises linearly at rate } b\text{)}.$$

Deriving the solution to the utility maximization problem given in (3) under a tax τ or standard $\bar{x}$ is straightforward. We consider policies that achieve a fixed total provision of $\widetilde{X}$ and bind on all agents. Specifically, we set $\widetilde{X}$ such that $\bar{x} \geq \hat{x}_g$ (as in proposition 3). Defining Δ_{ts} as the cost advantage that a standard has over a price-based policy, we obtain the following simple expression:

19. The intuition is that a year's change in one country will affect global climate only slightly, while any nonlinearity in benefits is likely to appear only for much larger temperature movements.

$$\Delta_{ts} = \alpha(1-\alpha)\frac{m^2(\widetilde{V}-1)^2}{b}. \quad (6)$$

First notice that the advantage of the standard is increasing in the square of both m and $\widetilde{V}$. The greater the strength of prosocial preferences ($\widetilde{V}$) or the marginal benefits from the public good (m), the bigger is the initial wedge in choice of provision across the two types. The standard overcomes this difference, achieving the first-best, whereas the price-based policy maintains this wedge and is thus a more costly way to attain $\widetilde{X}$.

Next we observe that as b declines and marginal costs become flatter, the advantage of the standard increases even further. This is at first counterintuitive, but notice that the levels of $\widetilde{V}$ and m alone determine the absolute difference in marginal costs between types. Holding the difference in marginal costs fixed, small values of b imply large differences in absolute levels of private provision.[20]

We next consider how the share of prosocial agents, α, affects the relative advantage of a standard. If everyone is prosocial or no one is ($\alpha = 1$ or 0), our model reduces to the standard equivalence between the two policies. Proposition 5 shows that the advantage of the standard has a single peak in α. In this example the peak occurs at $\alpha = ½$. Intuitively this is where the degree of heterogeneity in the population is maximized and the benefit of reallocating provision across agents is the greatest.

We finally explore the case in which heterogeneity exists in both prosocial preferences and marginal costs of provision. The effect of cost heterogeneity on optimal policy choice is quite intuitive in this setting. Allowing marginal costs to differ, $h'(x_g) = b_g x_g$ and $h'(x_u) = b_u x_u$, and solving as before yields

$$\Delta_{ts} > 0 \Leftrightarrow m(\widetilde{V}-1) > |b_g - b_u|\widetilde{X}. \quad (7)$$

That is, the standard is preferred as long as the wedge between the high and low valuations is greater than the absolute difference in agents' marginal costs of provision. When preference heterogeneity is relatively large, the standard dominates. When cost heterogeneity is relatively large, the price-based policy dominates.

3. POLICY DISCUSSION

The trade-off between preference heterogeneity across agents (favoring a standard) and heterogeneity in costs of provision (favoring a price instrument) appears in a range of policy directed at private provision of public goods. In our setting, a price instrument subsidizes provision of the public good uniformly for all agents. With respect to local public goods, a homeowner association (HOA) subsidizing lawn maintenance

20. We cannot examine the case as b goes to zero in this parameterization since the prosocial agent's private provision, $\hat{x}_g$, tends to infinity, removing the need for policy.

expenses is a uniform price instrument. Conversely, some minimum level of mandated maintenance represents our minimum standard.

This section provides discussion of three policy environments that demonstrate the application of our model: the first is a neighborhood amenities case, with incentives to provide public goods such as landscaping and sidewalk repair. This is a setting in which HOA and city-level policies are very common and both types of policy (prices and mandates) appear in practice.

We next turn to two energy-efficiency examples. The role for cost heterogeneity in these cases is much greater, favoring a price-based policy, and we discuss how the different types of heterogeneity might compete. We limit ourselves to examples in which the public good (energy conservation in these cases) comes from a single margin: durable good choice. When there is an important utilization margin after the durable is purchased, the benefits of a policy that can influence both margins at the same time (such as a fuel tax) become first-order relative to the distortions from heterogeneity that we model.

3.1. Neighborhood Amenities: Landscaping

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. The scale of the public good in these cases is such that every household can plausibly affect the aggregate level, and different households may have very different tastes. Improving the appearance of a home and lawn is costly in terms of both time and money. It is also reasonable that the marginal costs of improving the appearance of a home are upward sloping: mowing the lawn and trimming are relatively low cost whereas manicuring and fostering an extensive garden are more costly. Both, though, contribute to improving the visual aesthetics of the neighborhood. Finally, homeowners have varying degrees of affinity toward the appearance of lawns in their neighborhood. This is clearest when a homeowner with a strong preference for manicured lawns lives next door to another who is unconcerned with overgrown yards.

Neighborhoods often form HOAs to coordinate activities to improve neighborhood public goods such as landscaping. In lieu of an efficient but impractical set of individual-specific subsidies, the HOA has at least two second-best policy options to increase public good provision, mirroring the policy choices we model. First, they could subsidize homeowner spending on landscape services. This policy is akin to offering a uniform (across agents) subsidy for the public good in our model: households will choose different levels of the public good depending on their preferences and receive more subsidy if they provide more of the public good. Second, the HOA could mandate that all homes purchase some common minimum level of service from a landscaper. In this case some households could still provide more than the minimum level (as in proposition 4 above) but would not receive any additional subsidy or rebate.[21]

Our model suggests that the minimum standard could be more efficient than the subsidy incentive if households have heterogeneous preferences for neighborhood amenities such as landscaping. A subsidy provided for landscaping will increase provision by households who already have a preference for it, causing them to move farther out the marginal cost curve. The subsidy would be better spent if it could be targeted to those providing the lowest level (where marginal improvements are cheapest). Conversely, the policy with mandated basic lawn care will enforce public good provision by households with no preference for it, taking advantage of these low-cost improvements.

The key countervailing force in our model, arguing for a subsidy instead of the standard, comes from cost heterogeneity. In this particular context, cost heterogeneity could come simply from differing values of time spent gardening or, more subtly, through differing private valuation of one's own landscaping that will be correlated with the degree of public contribution. There is surely at least some amount of cost heterogeneity faced by individual agents in this setting, and the degree (relative to preference heterogeneity for the public good in aggregate) will determine which of the policies is optimal. In practice this is an empirical question and will likely differ across HOAs; we note that in practice different associations have different policies, and both minimum standards and subsidies are represented.

3.2. Refrigeration

Refrigeration technologies provide a relatively simple example of the setting we have in mind for durable goods, subject to prosocial "green" agents being motivated by Charness and Rabin (2002) efficiency preferences. In this case, households that want to provide a public good (energy conservation) through the purchase of more efficient refrigerators face rising marginal costs as reflected in the price of ever more efficient models. Utilization is nearly fixed since refrigeration is effective in only a narrow temperature range, and marginal costs are close to homogeneous across households if they face a common market price for appliances. Assuming that agents are motivated to provide energy conservation via a preference for social efficiency, differing preferences over large-scale environmental externalities get magnified, and we will see different contribution levels even when agents have the same costs. If green agents are motivated to

21. Whether the minimum mandated service is paid for by HOA fees or households directly is a distributional issue; the key distinction in our model is that provision below the minimum is not possible and provision above the minimum does not receive any (additional) subsidy.

provide the public good for other reasons, for example, private utility gains due to warm glow, then cost heterogeneity is the main factor at play.

In our setting a uniform price instrument subsidizes more efficient durables throughout the range of energy efficiency and does not discriminate across households. Consider a household's decision to purchase refrigerators with different embedded levels of pollution (where the private cost faced is the total price in present value of owning and operating the appliance). If the household considers the public good aspect of their choice (air pollution from electricity usage), this appears to the consumer as energy-efficient refrigerators being subsidized and becoming cheaper relative to less efficient ones. Alternatively, a minimum standard on the efficiency of refrigerators will appear as a narrowing of the choice set: it is no longer possible to buy a refrigerator that pollutes more than a certain amount. Price-based incentives therefore move all consumers (both green and nongreen) toward more efficient refrigerators than they would choose in the absence of the policy. This preserves the marginal cost difference between greens and nongreens, moving greens farther out a rising marginal cost curve while failing to take advantage of lower-cost improvements available to nongreens. A minimum standard could bring the two types closer together.

This system again has cost heterogeneity present as well: the degree (and correlation if present) of the two types of heterogeneity will determine optimal policy. In this case cost heterogeneity can appear in many ways, for example, as differences in electricity price (from increasing block rates or geographical differences), differences in discount rates changing the present value calculation, individualistic warm glow, and differing tastes for new technologies. Without the magnification of preference heterogeneity embedded in efficiency preferences (see app. C), these cost differences almost surely dominate, making a price instrument the best policy. We argue that additional empirical work is needed to understand the motivations for green purchases and the possibility that preference heterogeneity on the public goods side could also be at work. A further caveat to applying our model comes through the second margin of choice implicit in refrigerator size. Standards will not incentivize the choice of a smaller refrigerator, missing a potentially important margin for reducing energy use; a more complete model assessing the elasticity along this, and any additional margins of choice, would be needed.

3.3. Lighting Technology Choice

Lightbulb choice provides an example demonstrating the role of both cost and preference heterogeneity. As with refrigerators, the public goods framework applied to lighting only if green agents exhibit Charness and Rabin (2002) preferences. In addition, lighting could have an important use dimension: if lighting costs decrease, it is possible that households light their homes more intensely. We highlight below how this use dimension mitigates the applicability of our model to the lighting example.

Each household works its way out an increasing marginal cost curve by selecting one of a variety of ever more expensive bulbs: standard incandescents, halogen incandescents, CFLs, and LED-based bulbs. Consumers differ widely in their choices, spreading out across the spectrum of bulb options. Part of this difference may be attributable to green preference heterogeneity of the sort we model here and part of it to heterogeneity in the utility cost of things such as the color and flicker of light from different bulbs. We argue that the two forms of heterogeneity compete: in cases in which the bulbs are almost identical on aspects other than the provision of energy conservation (standard vs. halogen incandescents may come close to this case), minimum standards would dominate. In cases in which utility costs of different bulbs vary widely across consumers (likely relevant in the next switch from halogen incandescents to CFLs), the price-based incentives will again dominate.

Lightbulb choice is also the focus of several federal and state policies in the United States. "Price-based" policies in this case are simply subsidies to more efficient bulbs and are historically the predominant way to stimulate the purchase of energy-efficient lighting.[22] The alternative, minimum, standards have been introduced more recently. For example, California's Assembly Bill 1109 places a minimum standard that phased out standard incandescents between 2011 and 2013. Under this legislation, halogen incandescents became the minimum-efficiency bulb permitted by the standard. A federal minimum standard, roughly 1 year behind the California law in timing, has similar provisions. Yet the federal law proved contentious to the point that Congress acted to delay enforcement in late 2011.[23]

Our model provides a way of considering the trade-offs relevant in this choice of policy: the more conventional policy of subsidizing efficient bulbs moves everyone's choice farther out the technology cost curve, preserving a gap between greens and nongreens. The minimum standards instead push all consumers only as far as the halogen bulbs, taking advantage of a low-cost conservation option available to nongreens without distorting the incentives faced by greens. As we argue above, such a policy may be able to provide the public good (energy conservation) at lower cost: the greens purchasing CFLs or LED bulbs without policy will presumably continue to do so but will not be pushed any farther out the cost curve. This situation again corresponds to that considered in proposition 4, where the standard is binding on only one type of consumer.

Cost heterogeneity again enters, in the same ways as the examples above, and now perhaps through additional factors such as aversion to flicker or certain colors of light, ability to work with dimmers, bulb shape, and so on. Many of these issues become

---

22. "The New Light Bulbs Lose a Little Shine," *Wall Street Journal* (January 19, 2011), describes $548 million paid in subsidies to CFLs in California.

23. See "Let There Be Light Bulbs," *Wall Street Journal* (July 15, 2011), and the temporary suspension of the law, "Congress Kills Light Bulb Ban—Sort Of," *Forbes* (December 16, 2011).

even more relevant if the standard gets stricter, phasing out the halogen incandescents and allowing only CFL or LED lighting.[24] The greater this sort of heterogeneity, the greater the likelihood that the price-based instrument and not the minimum standard will dominate.

A caveat here is that our model considers only one margin of choice. We can compare competing policy interventions for lightbulb choice because they act on exactly the same margin. However, if households change their lighting intensity as a function of electricity cost, this creates two different margins that could provide the public good. Options then might include two policies (one for each margin) or a single policy such as a tax that can influence both margins at once. We cannot draw comparisons with other policies in the portfolio that reduce the utilization of lighting, such as normative messages that promote turning off lights when leaving a room or similar behavioral adjustments (e.g., Allcott 2011; Costa and Kahn 2013; Ferraro and Price 2013; Ito et al. 2013; Metcalfe and Dolan 2013), without a more complex model to allow interactions between green preferences and private costs along the two margins together.

4. CONCLUSION

We model an economy populated by two types of agents that differ with respect to the benefits they receive from a public good. If both types of agents face identical, but convex, costs of provision, the private equilibrium involves high-valuation agents working farther out their marginal cost curve and providing units of the public good that are more costly on the margin than those provided by low-valuation agents. We use this basic framework to compare the relative performance of different second-best policies designed to promote increased private provision of public goods. We show that minimum standards can provide an increase in aggregate provision of a public good at lower cost than price-based incentives if consumers face identical marginal cost curves. Intuitively, such standards tend to reduce the differences in marginal cost of provision across types and therefore reduce total cost. Price-based policies, on the other hand, place the same incentive on all individuals and so preserve inefficiency stemming from uneven provision of the public good.

Preferences are not the only dimension along which heterogeneity is likely to arise, however: in most applied settings agents also face heterogeneous costs of provision. We show how the two sources of heterogeneity compete in determining the preferred second-best policy. When cost heterogeneity is relatively large in magnitude a price-based policy dominates. The opposite holds when heterogeneity in prosocial preferences is the more important component. Whether preference or cost heterogeneity is relatively more or less important is an empirical question that must be answered in any

24. Relative to incandescents, the expected life span of CFLs is very different. Similarly, issues of color, warm-up times, and flicker may be important to some consumers but not to others.

given policy setting. Our results suggest that optimal policy choice can depend pivotally on the degree and correlation of heterogeneity along these two dimensions.

Several extensions of the model could lead to applications in more complex policy scenarios: Perhaps most important, the intensity of utilization, for example, distance traveled in vehicles, plays a critical role in total energy use.[25] In such settings, one would need to consider a broader set of policies and additional dimensions of preference heterogeneity. Two policies may be necessary in the first-best, and effectiveness along the two margins simultaneously becomes a critical test for second-best policies. Further, we employ a model with only two types of agents while a continuum of preferences might better reflect empirical differences in the provision of public goods. Crowding out, both within and between types of agents, also reflects an important possibility that could be incorporated by a more detailed model.

Finally, our model also raises an interesting dilemma concerning incidence: The cost-minimizing standard places greater burden on those who value the public good least. An interesting direction for future work is thus to explore the distributional impacts of various policy options and how the costs of providing a given level of the public good would change with the imposition of a compensation requirement as in Bovenberg et al. (2008). Careful analysis of incidence could also contribute to the larger political economy question of how the form of heterogeneity in preferences for a public good, and policy designed with this in mind, affects optimal public good levels.

APPENDIX A

Proof of Lemma 1

By definition, $\widetilde{V}$ is strictly larger for those with high valuation. Private equilibrium provision for high-valuation agents, $\hat{x}_g$, and low-valuation agents, $\hat{x}_u$, is implicitly defined by the following system of equations:

$$h'(\hat{x}_g) = \widetilde{V} f'(\alpha N \hat{x}_g + (1 - \alpha) N \hat{x}_u), \tag{A1}$$

$$h'(\hat{x}_u) = f'(\alpha N \hat{x}_g + (1 - \alpha) N \hat{x}_u). \tag{A2}$$

Since $\widetilde{V} > 1$, it must be that $\hat{x}_g > \hat{x}_u$, giving the desired result.

Proof of Proposition 1

Set up the cost minimization problem directly with a Lagrangian such that

$$\min_{\underline{x}} \sum_i h(x_i) \tag{A3}$$

25. Hausman (1979) and a rich subsequent literature consider interactions between purchase and utilization of durables, e.g.

$$\text{s.t.} \sum_i x_i = \widetilde{X} \tag{A4}$$

$$\rightarrow L = \sum_i h(x_i) + \lambda\left[\widetilde{X} - \sum_i x_i\right]. \tag{A5}$$

The first-order conditions for this Lagrangian are $h'(x_i) = \lambda$ for all $i = 1, 2, \ldots, N$, implying that $h'(x_i) = h'(x_j)$ for all i, j, which gives the desired result.

Proof of Corollary 1

Economywide social efficiency preferences are defined by parameters $\alpha = 1$ and $e = 1 \rightarrow \widetilde{V} = N$ (see fn. 10). In this case, the privately provided equilibrium is defined by the single equation $h'(\hat{x}_g) = Nf'(N\hat{x}_g)$. Accounting only for the direct benefits accruing to agents (e.g., not the social efficiency component of their preference), the socially optimal level of the public good satisfies the following condition:

$$\sum_i h'(x_i^*) = Nf'(Nx_i^*). \tag{A6}$$

Equation (A6) states that the sum of the marginal costs for each agent's provision must equal the sum of the marginal benefits. Summing up the direct effect of all agents' private provision gives $\sum h'(\hat{x}_g) = Nf'(N\hat{x}_g)$. By strict concavity of $h(\cdot)$, $x_g = x_i^*$, giving the desired result.

Proof of Proposition 2

There are four conditions that must be jointly satisfied in order to have efficient public good provision with an asymmetric price instrument profile. By proposition 1 and proposition 3, these conditions are

$$\begin{aligned} h'(x_g^t) &= \tau_g + \widetilde{V}f'(\widetilde{X}), \\ h'(x_u^t) &= \tau_u + f'(\widetilde{X}), \\ \widetilde{X} &= \alpha N x_g^t + (1-\alpha)N x_u^t, \\ h'(x_g^t) &= h'(x_u^t). \end{aligned}$$

Substituting in, we get the condition $\widetilde{V}f'(\widetilde{X}) + \tau_g = f'(\widetilde{X}) + \tau_u$. Rearranging gives $(V-1)f'(\widetilde{X}) = \tau_u - \tau_g$. By assumption, $V > 1$ and $f'(\cdot) > 0$, giving the desired result.

Proof of Proposition 3

Choose a price instrument τ such that $\sum_i x_i = \widetilde{X}$. The price instrument, τ, enters the budget constraint as $y = c + h(x_i) - \tau x_i$. Agents still privately optimize such that private equilibrium is jointly determined by

$$h'(\hat{x}_g^t) = \tau + \widetilde{V}f'(\widetilde{X}), \tag{A7}$$

$$h'(\hat{x}_u^t) = \tau + f'(\widetilde{X}). \tag{A8}$$

By convexity of $h(\cdot)$, $h'(\hat{x}_g^t) \neq h'(\hat{x}_u^t)$. Under a standard, all agents provide a minimum level of provision $\widetilde{X}/N$ such that $h'(\widetilde{X}/N) > \widetilde{V}f'(\widetilde{X}) > f'(\widetilde{X})$. As a result, $\widetilde{X}$ is provided such that $h'(x_i) = h'(x_j)$ for all i, j. A standard is a least-cost mechanism for providing $\widetilde{X}$ whereas by equations (A7) and (A8) a tax is not, giving the desired result.

Proof of Proposition 4

Equilibrium in the price instrument case is given by the system

$$h'(\hat{x}_g^t) = \tau + \widetilde{V}f'(\widetilde{X}), \tag{A9}$$

$$h'(\hat{x}_u^t) = \tau + f'(\widetilde{X}). \tag{A10}$$

By convexity of $h(\cdot)$, $h'(\hat{x}_g^t) > h'(\hat{x}_u^t)$ and $x_g^t > x_u^t$. In the case of the standard, prosocial agents' provision, x_g^s, is defined by their first-order condition: $h'(\hat{x}_g^s) = \widetilde{V}f'(\widetilde{X})$. By convexity of $h(\cdot)$ and $\tau > 0$, it implies $x_g^t > x_g^s$. Further, $\widetilde{X} = N(\alpha x_g^s + (1-\alpha)x_u^s)$. As a result, low-valuation agents' provision of the public good in the case of standards can be expressed as $x_u^s = ((\widetilde{X}/N) - \alpha x_g^s)/(1-\alpha)$. Since $x_g^t > x_g^s$, it implies $x_u^s > x_u^t$ and subsequently $x_g^t - x_u^t > x_g^s - x_u^s$.

Consider a case in which $x_g^t = x_g^s + \epsilon$ for $\epsilon > 0$ to reach some $\widetilde{X}$. Note that the distribution of provision in the case of a standard would therefore be

$$\alpha N x_g^s + (1-\alpha)N\left(x_u^t + \frac{\alpha}{1-\alpha}\epsilon\right).$$

The average cost of provision for x across agents in the price instrument case is $p^t = \alpha h(x_g^s + \epsilon) + (1-\alpha)h(x_u^t)$ and in the case of subsidies is therefore

$$p^s = \alpha h(x_g^s) + (1-\alpha)h\left(x_u^t + \frac{\alpha}{1-\alpha}\epsilon\right).$$

By concavity of $h(\cdot)$, $p^t > p^s$, giving the desired result.

Proof of Corollary 2

Equilibrium in the case of price instruments is given by equations (A7) and (A8) above in addition to the level constraint: $\widetilde{X} = \alpha N x_g^t + (1-\alpha)N x_u^t$. Substituting a transformation of the constraint in for low-valuation agents' provision, $x_u^t = ((\widetilde{X}/N) - \alpha x_g^t)/(1-\alpha)$, leaves two equations and two unknowns. Cramer's rule states that

$$\frac{dx_g^t}{d\alpha} = \frac{|\Lambda_{1,\alpha}|}{|H|}, \quad \frac{d\tau}{d\alpha} = \frac{|\Lambda_{2,\tau}|}{|H|}, \tag{A11}$$

where H is the Hessian of the system and $\Lambda_{n,\phi}$ is the Hessian with the nth column replaced with the negatives of the first-order condition derivatives with respect to the parameter ϕ:

$$H = \begin{pmatrix} \frac{\partial FOC_g}{\partial x_g^t} & \frac{\partial FOC_g}{\partial \tau} \\ \frac{\partial FOC_u}{\partial x_g^t} & \frac{\partial FOC_u}{\partial \tau} \end{pmatrix} = \begin{pmatrix} -h''(x_g^t) & 1 \\ h''(x_g^t)\frac{\alpha}{1-\alpha} & 1 \end{pmatrix}. \tag{A12}$$

By inspection, the determinant of the Hessian H is negative. Further, $\Lambda_{1,\alpha}$ and $\Lambda_{2,\tau}$ are, respectively,

$$\begin{aligned} \Lambda_{1,\alpha} &= \begin{pmatrix} 0 & 1 \\ h''(x_g^t)\frac{(\widetilde{X}/N) - x_g^t}{(1-\alpha)^2} & 1 \end{pmatrix}, \\ \Lambda_{2,\tau} &= \begin{pmatrix} -h''(x_g^t) & 0 \\ h''(x_g^t)\frac{\alpha}{1-\alpha} & h''(x_g^t)\frac{(\widetilde{X}/N) - x_g^t}{(1-\alpha)^2} \end{pmatrix}. \end{aligned} \tag{A13}$$

Noting that $\widetilde{X}/N < x_g^t$ by proposition 5, by inspection $|\Lambda_{1,\alpha}| > 0$ and $|\Lambda_{2,\tau}| > 0$, implying that $dx_g^t/d\alpha < 0$ and $d\tau/d\alpha < 0$, giving the desired result.

Proof of Corollary 3

Equilibrium in the case of price instruments is given by equations (A7) and (A8) above in addition to the level constraint: $\widetilde{X} = \alpha N x_g^t + (1-\alpha)N x_u^t$. Substituting a transformation of the constraint in for low-valuation agents' provision, $x_u^t = ((\widetilde{X}/N) - \alpha x_g^t)/(1-\alpha)$, leaves two equations and two unknowns. Cramer's rule states that

$$\frac{dx_g^t}{d\widetilde{V}} = \frac{|\Lambda_{1,\tilde{V}}|}{|H|}, \quad \frac{d\tau}{d\widetilde{V}} = \frac{|\Lambda_{2,\tilde{V}}|}{|H|}, \tag{A14}$$

where H is the Hessian of the system and $\Lambda_{n,\phi}$ is the Hessian with the nth column replaced with the negatives of the first-order condition derivatives with respect to the parameter ϕ. By corollary 2, $|H| < 0$. Further $|\Lambda_{1,\tilde{V}}|$ and $|\Lambda_{2,\tilde{V}}|$ are defined as

$$\begin{aligned} \Lambda_{1,\tilde{V}} &= \begin{pmatrix} -f'(\widetilde{X}) & 1 \\ 0 & 1 \end{pmatrix}, \\ \Lambda_{2,\tilde{V}} &= \begin{pmatrix} -h''(x_g^t) & -f'(\widetilde{X}) \\ h''(x_g^t)\frac{\alpha}{1-\alpha} & 0 \end{pmatrix}. \end{aligned} \tag{A15}$$

By inspection, $|\Lambda_{1,\tilde{V}}| < 0$ and $|\Lambda_{2,\tilde{V}}| > 0$. Therefore, $dx_g^t/d\widetilde{V} > 0$ and $d\tau/d\widetilde{V} < 0$, giving the desired result.

Proof of Proposition 5

This proof proceeds by construction. We first show that the limit of the total derivative of the difference, Δ_{ts}, is positive as $\alpha \to 0^+$ and negative as $\alpha \to 1^-$. We then show that the partial derivative of the difference between the policies, Δ_{ts}, is positive. We can define Δ_{ts} as

$$\Delta_{ts} = \alpha N\int_{\tilde{X}/N}^{x_g^t} h'(x)dx - (1-\alpha)N\int_{x_u^t=((\tilde{X}/N)-\alpha N)/(1-\alpha)}^{\tilde{X}/N} h'(x)dx. \tag{A16}$$

The total derivative of equation (A16) can be found using Leibniz's rule:

$$\begin{aligned}\frac{d\Delta_{ts}}{d\alpha} = {} & N\int_{\tilde{X}/N}^{x_g^t} h'(x)dx + \alpha N\left(Mh'\left(x_g^t\right)\right) + N\int_{x_u^t=((\tilde{X}/N)-\alpha N)/(1-\alpha)}^{\tilde{X}/N} h'(x)dx \\ & - (1-\alpha)N\left(-M\frac{\left(\widetilde{X}/N\right)-x_g^t}{(1-\alpha)^2}h'\left(\frac{\left(\widetilde{X}/N\right)-\alpha x_g^t}{1-\alpha}\right)\right),\end{aligned} \tag{A17}$$

$$M \equiv \frac{\left(h''(x_u^t)\dfrac{\widetilde{X}}{N} - x_g^t\right)/(1-\alpha)^2}{h''(x_g^t) + \dfrac{\alpha}{1-\alpha}h''(x_u^t)}. \tag{A18}$$

Note that $M < 0$ for any α and consider $\lim_{\alpha\to 0^+}$. The terms with the integrals converge to zero, leaving only the terms multiplying M. The first term goes to zero and, as above, the second term is positive such that $\lim_{\alpha\to 0^+} > 0$. Similarly, $\lim_{\alpha\to 1^-}$ is signed by the first term multiplying M, which is positive, so $\lim_{\alpha\to 1^-} < 0$. Finally, the partial derivative of Δ_{ts} is

$$\begin{aligned}\frac{\partial\Delta_{ts}}{\partial\alpha} = {} & N\int_{\tilde{X}/N}^{x_g^t} h'(x)dx + N\int_{x_u^t=((\tilde{X}/N)-\alpha N)/(1-\alpha)}^{\tilde{X}/N} h'(x)dx \\ & - (1-\alpha)\,N\left(-M\frac{(\widetilde{X}/N)-x_g^t}{(1-\alpha)^2}h'\left(\frac{(\widetilde{X}/N)-\alpha x_g^t}{1-\alpha}\right)\right).\end{aligned} \tag{A19}$$

By inspection, equation (A19) is positive, completing the proof.

Proof of Corollary 4

This proof proceeds by construction. We show that the total derivative of the difference, Δ_{ts}, is everywhere greater than zero. Again using Leibniz's rule and simplifying, we find

$$\frac{d\Delta_{ts}}{d\widetilde{V}} = \alpha N \frac{dx_g^t}{d\widetilde{V}}\left(h'(x_g^t) - h'\left(\frac{(\widetilde{X}/N) - \alpha N}{1-\alpha}\right)\right). \tag{A20}$$

By proposition 5, corollary 3, and convexity of $h(\cdot)$, equation (A20) is positive, giving the desired result.

Proof of Proposition 6

This proof proceeds by construction and is similar to that of corollary 2. Using the same notation as in corollary 2,

$$\frac{dx_g^t}{d\delta} = \frac{|\Lambda_{1,\alpha}|}{|H|}, \quad \frac{d\tau}{d\delta} = \frac{|\Lambda_{2,\tau}|}{|H|}. \tag{A21}$$

These matrices are defined as

$$H = \begin{pmatrix} \dfrac{\partial FOC_g}{\partial x_g^t} & \dfrac{\partial FOC_g}{\partial \tau} \\ \dfrac{\partial FOC_u}{\partial x_g^t} & \dfrac{\partial FOC_u}{\partial \tau} \end{pmatrix} = \begin{pmatrix} -h''(x_g^t)\delta & 1 \\ h''(x_g^t)\dfrac{\alpha}{1-\alpha} & 1 \end{pmatrix}. \tag{A22}$$

By inspection, the determinant of the Hessian H is negative. Further, $\Lambda_{1,\delta}$ and $\Lambda_{2,\delta}$ are, respectively,

$$\Lambda_{1,\delta} = \begin{pmatrix} -h'(x_g^t) & 1 \\ 0 & 1 \end{pmatrix},$$

$$\Lambda_{2,\delta} = \begin{pmatrix} -h''(x_g^t)\delta & -h'(x_g^t) \\ h''(x_g^t)\dfrac{\alpha}{1-\alpha} & 0 \end{pmatrix}. \tag{A23}$$

By inspection, the determinants of $\Lambda_{1,\delta}$ and $\Lambda_{2,\delta}$ are positive and negative, respectively. As a result, using Cramer's rule, $dx_g^t/d\delta < 0$ and $d\tau/d\delta > 0$, giving the desired result.

Proof of Proposition 7

The analogue of equation (A16) in this version of the model is

$$\Delta_{ts} = \alpha N(1+\delta)\int_{\tilde{X}/N}^{x_g^t} h'(x)dx - (1-\alpha)N\int_{x_u^t=((\tilde{X}/N)-\alpha x_g^t)/(1-\alpha)}^{\tilde{X}/N} h'(x)dx. \tag{A24}$$

By inspection, $\partial\Delta_{ts}/\partial\delta > 0$. However, accounting for indirect effects as in proposition 4, we take the total derivative using Leibniz's rule:

$$\frac{d\Delta_{ts}}{d\delta} = \alpha N \int_{\tilde{X}/N}^{x_g^t} h'(x)dx + \alpha N \frac{dx_g^t}{d\delta}\left(h'(x_g^t) - h'(x_u^t) + \delta h'(x_g^t)\right). \qquad \text{(A25)}$$

By inspection, the direct effect is still positive but the indirect effect operates in the other direction since $dx_g^t/d\delta < 0$ from proposition 6. The net effect, though, is still positive as long as

$$\frac{h(x_g^t) - h\left(\dfrac{\tilde{X}}{N}\right)}{-\dfrac{dx_g^t}{d\delta} h'(x_g^t)} > 1 + \delta. \qquad \text{(A26)}$$

This condition states that prosocial agents do not have a marginal cost of abatement curve that is so high they provide less of the public good than other agents. This will be satisfied by the negative correlation between costs and prosocial preferences. This concludes the proof.

APPENDIX B

This appendix shows equivalence between the model used in this paper and a model with a linear budget constraint but decreasing marginal utility with respect to the private good.

The utility and budget specification in this paper are represented as

$$U_i(c_i, X|\Theta) = c_i + \tilde{V} f_i(X) \quad \text{s.t. } y_i = c_i + h(x_i)$$
$$\Rightarrow U_i(c_i, X|\Theta) = y_i - h(x_i) + \tilde{V} f_i(X).$$

As shown above, the first-order condition of the consumer's problem is

$$h'(x_i^*) \geq \tilde{V} f'(X^*) \quad \forall i. \qquad \text{(B1)}$$

The function $h'(x_i^*)$ is the first derivative of a convex function.

It is possible to instead use a linear budget constraint with decreasing marginal utility of the numeraire consumption good. Assume that utility derived from the numeraire consumption good is described by a concave function $v(c)$, $v'(c) > 0$, and $v''(c) < 0$. The consumer's choice problem can then be expressed as

$$U_i(c_i, X|\Theta) = v(c_i) + \tilde{V} f_i(X) \quad \text{s.t. } y_i = c_i + p_x x_i$$
$$\Rightarrow U_i(c_i, X|\Theta) = v(y_i - p_x x_i) + \tilde{V} f_i(X). \qquad \text{(B2)}$$

Now consider the private equilibrium of the consumer given the model in equation (B2). The consumer's first-order condition is

$$-p_x v'(y_i - p_x x_i) \geq \tilde{V} f_i'(X), \qquad \text{(B3)}$$

with equality if $x_i > 0$. The left-hand side of the consumer's first-order condition in equation (B3) is the opportunity cost of spending additional resources on purchase of the public good. Specifically, $-p_x\nu'(y_i - p_x x_i)$ can be evaluated as a function of x_i in equilibrium as opposed to a function of c_i. If $\nu(c_i)$ is concave and increasing in c_i, then by definition it is concave and decreasing in x_i. Further, $-p_x\nu'(y_i - p_x x_i)$ is increasing in x_i. Figure B1 shows this relationship visually.

Importantly, the left-hand sides of equations (B1) and (B3) are both increasing functions of the arguments x_i. As a result, equilibrium in these models will be equivalent. For example, the proof of lemma 1 under this alternative specification is as follows:

Proof of lemma 1: By definition, $\widetilde{V}$ is strictly larger for high-valuation agents than for others. Private equilibrium provision for high-valuation agents, $\hat{x}_g$, and low-valuation agents, $\hat{x}_u$, is implicitly defined by the following system of equations:

$$-p_x\nu'(y_i - p_x\hat{x}_g) = \widetilde{V}f'(\alpha N\hat{x}_g + (1-\alpha)N\hat{x}_u), \tag{B4}$$

$$-p_x\nu'(y_i - p_x\hat{x}_u) = f'(\alpha N\hat{x}_g + (1-\alpha)N\hat{x}_u). \tag{B5}$$

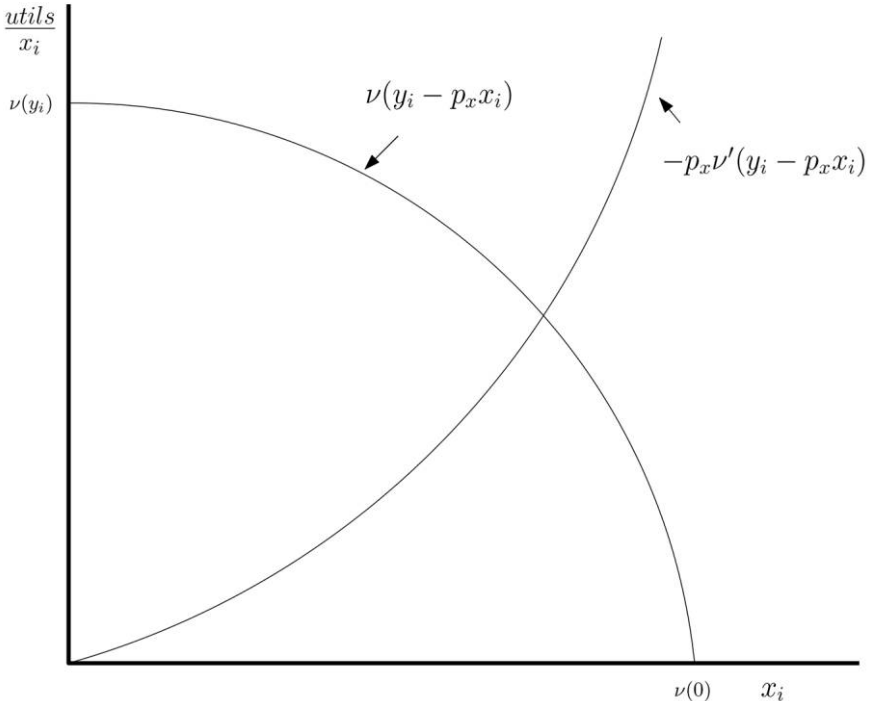

Figure B1. Equivalence of alternative utility specification

Since $\widetilde{V} > 1$, it must be that $\hat{x}_g > \hat{x}_u$ since $-p_x v'(y_i - p_x \hat{x}_g)$ is increasing in x_i, giving the desired result.

APPENDIX C

The goal of this appendix is to show how the public goods model we work with interacts with cases including a large population, for example, climate change and provision of energy efficiency by households. Begin by noting two distinct reasons an agent might contribute to a public good: The neoclassical reason elaborated in Bergstrom et al. (1986) is that they value, and can individually affect, the aggregate provision of the good. An alternative explanation is that the agent gets a private benefit from providing that accrues only if they personally do the provision (e.g., warm glow). We highlight in the paper how warm glow manifests as cost heterogeneity in our model: it lowers the cost to the agent of providing relative to other agents. This appendix relates to the neoclassical justification.

Imagine a public good like the one coming from household energy efficiency and the associated reduction in climate change impacts for the globe. Suppose for simplicity that global population is 7×10^9 and the total benefit from avoiding a ton of carbon emissions is constant at $40/ton. The benefit of provision to any individual agent ($f_i(X)$ in our notation) is clearly very small and is $5.7 \times 10^{-9} X$. If we add rising marginal costs, agents would provide only a tiny amount of conservation (such that their $MC = 5.7 \times 10^{-9}$). Now suppose some of the agents have a green preference that makes them like the public good 100 times more than the other agents. Even though they like the public good much more, the green agents are still very small. They get a marginal benefit of

$$(100 \times 40)/(7 \times 10^9) = 5.7 \times 10^{-7}.$$

The green agents would provide a tiny bit more of the public good (following their individual rising marginal cost curve farther out), so technically the effects from preference heterogeneity of the sort we model still appear. However, all agents (green or not) have benefits so close to zero that the effects we study would be insignificant relative to any form of cost heterogeneity. Our results therefore apply most generally in settings with a small N.

However, suppose the green agents have a preference for the public good along the lines of Charness and Rabin (2002):

$$V_i = f_i(X) + \sum_{j \neq i} f_j(X).$$

Agents with Charness and Rabin efficiency preferences effectively internalize the externalities associated with their own provision. The summation over all the other agents now makes the green preference quite strong. With these preferences green agents provide the public good until their marginal cost of provision is equal to

$$5.7 \times 10^{-9} + \sum_{j \neq i} 5.7 \times 10^{-9} = 40.$$

The summation inside the utility function means that N multiplies the green preference, so even as the agent's contribution becomes very small with large N, the value placed on aggregate provision increases to offset. This creates a substantial wedge between the greens and nongreens (the nongreens will provide only up to the near-zero marginal cost described above), and so the distortions from preference heterogeneity become important.

The key difference between this model and one of individual warm glow is that the green agents have $f_{j \neq i}(X)$ in their utility functions and so care about the aggregate level of X directly. This type of prosocial agent would receive utility even if others have provided the conservation. Since little empirical evidence is available on the form of underlying preferences, the most general application of our model will be to cases with a relatively small number of agents; but under certain forms of preference the heterogeneity can remain important even as N grows large.

REFERENCES

Akerlof, George, and Rachel Kranton. 2000. Economics and identity. *Quarterly Journal of Economics* 115 (3): 373–99.

Allcott, Hunt. 2011. Social norms and energy conservation. *Journal of Public Economics* 95 (9): 1082–95.

Allcott, Hunt, Sendhil Mullainathan, and Dmitry Taubinksy. 2012. Externalities, internalities and the targeting of energy policy. Working paper, New York University.

Allcott, Hunt, and Dmitry Taubinsky. 2015. Evaluating behaviorally-motivated policy: Experimental evidence from the lightbulb market. *American Economic Review* 105 (8): 2501–38.

Andreoni, James. 1989. Giving with impure altruism: Applications to charity and Ricardian equivalence. *Journal of Political Economy* 97 (6): 1447–58.

Andreoni, James, Justin Rao, and Hannah Trachtman. 2017. Avoiding the ask: A field experiment on altruism, empathy, and charitable giving. *Journal of Political Economy* (forthcoming).

Becker, Gary. 1974. A theory of social interactions. *Journal of Political Economy* 82 (6): 1063–93.

Benabou, Ronald, and Jean Tirole. 2006. Incentives and prosocial behavior. *American Economic Review* 96 (5): 1652–78.

Bergstrom, Ted, Lawrence Blume, and Hal Varian. 1986. On the private provision of public goods. *Journal of Public Economics* 29 (1): 25–49.

Bovenberg, Lans, Lawrence Goulder, and Mark Jacobsen. 2008. Costs of alternative environmental policy instruments in the presence of industry compensation requirements. *Journal of Public Economics* 92 (5): 1236–53.

Brandts, Jordi, and Arthur Schram. 2001. Cooperation and noise in public goods experiments: Applying the contribution function approach. *Journal of Public Economics* 79 (2): 399–427.

Carson, Richard, and Jacob LaRiviere. 2017. Structural uncertainty and pollution control: Optimal stringency with unknown pollution sources. *Environmental and Resource Economics* (forthcoming).

Charness, Gary, and Matthew Rabin. 2002. Understanding social preferences with simple tests. *Quarterly Journal of Economics* 117 (3): 817–69.

Costa, Dora, and Matthew Kahn. 2013. Energy conservation nudges and environmentalist ideology: Evidence from a randomized residential electricity field experiment. *Journal of the European Economic Association* 11 (3): 680–702.

Dana, Jason, Daylian Cain, and Robyn Dawes. 2006. What you don't know won't hurt me: Costly (but quiet) exit in dictator games. *Organizational Behavior and Human Decision Processes* 100 (2): 193–201.

Dana, Jason, Roberto Weber, and Jason Kuang. 2007. Exploiting moral wiggle room: Experiments demonstrating an illusory preference for fairness. *Economic Theory* 33 (1): 67–80.

DellaVigna, Stefano, John List, and Ulrike Malmendier. 2012. Testing for altruism and social pressure in charitable giving. *Quarterly Journal of Economics* 127 (1): 1–56.

DellaVigna, Stefano, John List, Ulrike Malmendier, and Gautam Rao. 2013. The importance of being marginal: Gender differences in generosity. *American Economic Review* 103 (3): 586–90.

Diamond, Peter. 1973. Consumption externalities and imperfect corrective pricing. *Bell Journal of Economics* 4 (2): 526–38.

Ferraro, Paul, and Michael Price. 2013. Using non-pecuniary strategies to influence behavior: Evidence from a large scale field experiment. *Review of Economics and Statistics* 95 (1): 64–73.

Fischbacher, Urs, and Simon Gaechter. 2008. Heterogeneous social preferences and the dynamics of free riding in public goods experiments. TWI Research Paper Series, 27, Thurgauer Wirtschaftsinstitut, Universität Konstanz.

Fowlie, Meredith, and Nicholas Muller. 2013. Market-based emissions regulation when damages vary across sources: What are the gains from differentiation? NBER Working paper 18801, National Bureau of Economic Research, Cambridge, MA.

Glazer, Amihai, and Kai Konrad. 1996. A signaling explanation for charity. *American Economic Review* 86 (4): 1019–28.

Hausman, Jerry. 1979. Individual discount rates and the purchase and utilization of energy-usage durables. *Bell Journal of Economics* 10 (1): 387–89.

Hoel, Michael, and Larry Karp. 2001. Taxes and quotas for a stock pollutant with multiplicative uncertainty. *Journal of Public Economics* 82:91–114.

Holland, Stephen. 2012. Taxes and trading versus intensity standards: Second-best environmental policies with incomplete regulation (leakage) or market power. *Journal of Environmental Economics and Management* 63 (3): 375–87.

Hollander, Heinz. 1990. A social exchange approach to voluntary cooperation. *American Economic Review* 80 (5): 1157–67.

Ito, Koichiro, Takanori Ida, and Makoto Tanaka. 2013. The persistence of moral suasion and economic incentives: Field experimental evidence from energy demand. NBER Working paper 20910, National Bureau of Economic Research, Cambridge, MA.

Jacobsen, Grant, Matt Kotchen, and Michael Vandenbergh. 2012. The behavioral response to voluntary provision of an environmental public good: Evidence from residential electricity demand. *European Economic Review* 56 (5): 946–60.

Jacobsen, Mark. 2013. Evaluating US fuel economy standards in a model with producer and household heterogeneity. *American Economic Journal: Economic Policy* 5 (2): 148–87.

Kessler, Judd, and Katherine Milkman. 2014. Identity and charitable giving. Working paper, University of Pennsylvania.

Kotchen, Matt, and M. R. Moore. 2007. Private provision of environmental public goods: Household participation in green electricity programs. *Journal of Environmental Economics and Management* 53:1–16.

Krupka, Erin, and Roberto Weber. 2013. Identifying social norms using coordination games: Why does dictator game sharing vary? *Journal of the European Economic Association* 11 (3): 495–524.

Kurzban, Robert, and Daniel Houser. 2001. Individual differences in cooperation in a circular public goods game. *European Journal of Personality* 1 (3): S37–S52.

Lazear, Edward, Ulrike Malmendier, and Roberto Weber. 2012. Sorting in experiments with application to social preferences. *American Economic Journal: Applied Economics* 4 (1): 136–63.

Levine, David. 1998. Modeling altruism and spitefulness in experiments. *Review of Economic Dynamics* 1 (3): 593–622.

Martinot, Eric, and Nils Borg. 1998. Energy-efficient lighting programs: Experience and lessons from eight countries. *Energy Policy* 26 (14): 1071–81.

Metcalfe, Robert, and Paul Dolan. 2013. Neighbors, knowledge, and nuggets: Two natural field experiments on the role of incentives on energy conservation. Working paper, University of Chicago.

Poortinga, Wouter, Linda Steg, Charles Vlek, and Gerwin Wiersma. 2003. Environmental constraints on hydropower: An ex post benefit-cost analysis of dam relicensing in Michigan. *Journal of Economic Psychology* 24 (1): 49–64.

Saphores, Jean, Hilary Nixon, Oladele Gounseitan, and Andrew Shapiro. 2007. California households' willingness to pay for "green" electronics. *Journal of Environmental Planning and Management* 50 (1): 113–33.

Segerson, Kathleen. 1988. Uncertainty and incentives for nonpoint pollution control. *Journal of Environmental Economics and Management* 15 (1): 87–98.

Tsvetanov, Tsvetan, and Kathleen Segerson. 2014. The welfare effects of energy efficiency standards when choice sets matter. *Journal of the Association of Environmental and Resource Economists* 1 (1): 233–71.

Weitzman, Martin. 1974. Prices versus quantities. *Review of Economic Studies* 41:477–91.