EconS 501 – Micro Theory I
Recitation #13 – Public Goods[1]

1. **Public goods with different degrees of publicness.** Two consumers ($i = 1, 2$) each have income M to split between two goods. *Good 1* delivers one unit of consumption to its purchaser and $\gamma \in [0, 1]$ units to the other consumer. Each consumer i has utility

$$U^i = \ln(x_1^i) + x_2^i,$$

where x_k^i is consumer i's consumption of good k.

   (a) Provide an interpretation of γ.

   (b) Assume good 2 is a pure private good. Find the Nash equilibrium at prices $p_1 = p_2 = 1$.

   (c) Maximize the sum of utilities; show the equilibrium is Pareto-efficient iff $\gamma = 0$.

   (d) Now assume good 2 *also* has degree-of-publicness γ. Find the Nash equilibrium and show it is efficient for every γ.

   (e) Explain the conclusion in (d).

Answer. *Part (a): interpretation.* γ measures the "degree of publicness" of good 1. When $\gamma = 0$, good 1 is a pure private good (only the buyer consumes it). When $\gamma = 1$, good 1 is a pure public good (both consumers get one unit each from a single purchase). Intermediate values interpolate: an "impure" public good with partial spillovers.

Answer. *Part (b): Nash equilibrium (good 2 private).*

Step 1. Consumer 1's problem. Let y_1^i denote consumer i's *purchase* of good 1. Consumer 1's consumption of good 1 is $x_1^1 = y_1^1 + \gamma y_1^2$. With the budget $y_1^1 + x_2^1 = M$ (so $x_2^1 = M - y_1^1$) and taking y_1^2 as given (Nash),

$$U^1 = \ln(y_1^1 + \gamma y_1^2) + (M - y_1^1).$$

Step 2. FOC. Differentiate in y_1^1:

$$\frac{1}{y_1^1 + \gamma y_1^2} - 1 = 0 \iff y_1^1 + \gamma y_1^2 = 1.$$

Consumer 2's FOC is symmetric: $y_1^2 + \gamma y_1^1 = 1$.

[1]Felix Munoz-Garcia, School of Economic Sciences, Washington State University, 103H Hulbert Hall, Pullman, WA. Tel. 509-335-8402. Email: fmunoz@wsu.edu.

Step 3. Solve the symmetric equilibrium. In the symmetric Nash equilibrium $y_1^1 = y_1^2 \equiv y_1$, so

$$(1+\gamma)\,y_1 = 1 \iff y_1 = \frac{1}{1+\gamma}.$$

Each consumer's actual consumption of good 1 is $x_1 = (1+\gamma)\,y_1 = 1$.

Answer. *Part (c): social planner's problem and efficiency.*

Step 1. Social welfare function. Sum of utilities:

$$W = \left[\ln(y_1^1 + \gamma y_1^2) + M - y_1^1\right] + \left[\ln(y_1^2 + \gamma y_1^1) + M - y_1^2\right].$$

Step 2. Symmetric FOC. Applying the symmetry $y_1^1 = y_1^2 = y_1$ upfront, $W = 2\ln\big((1+\gamma)y_1\big) + 2(M - y_1)$. Differentiating,

$$\frac{2}{y_1} - 2 = 0 \iff y_1 = 1.$$

So the socially optimal purchase per consumer is $y_1^o = 1$, yielding consumption $x_1^o = (1+\gamma)\cdot 1 = 1+\gamma$.

Step 3. Compare. The Nash equilibrium had $y_1^m = 1/(1+\gamma)$; the social optimum has $y_1^o = 1$. These coincide iff $\gamma = 0$. For any $\gamma > 0$, the Nash equilibrium purchase falls short of the optimum: each consumer under-provides good 1 because she ignores the γ units of spillover benefit her purchase confers on the other consumer. Classic free-riding.

Answer. *Part (d): both goods with degree γ.*

Step 1. New utility. If good 2 also has publicness γ, consumer 1's utility becomes

$$U^1 = \ln(y_1^1 + \gamma y_1^2) + \left[(M - y_1^1) + \gamma(M - y_1^2)\right],$$

because consumer 1 now enjoys her own residual $M - y_1^1$ (spent on good 2) *plus* γ times consumer 2's residual.

Step 2. FOC and Nash equilibrium. Differentiating in y_1^1:

$$\frac{1}{y_1^1 + \gamma y_1^2} - 1 = 0 \iff y_1^1 + \gamma y_1^2 = 1.$$

(The γ term added by publicness of good 2 does not appear because we're differentiating in consumer 1's own choice, and $\gamma(M - y_1^2)$ doesn't depend on y_1^1.) Combined with the symmetric FOC $y_1^2 + \gamma y_1^1 = 1$, the equilibrium is again $y_1 = 1/(1+\gamma)$.

Step 3. Social planner's problem. With symmetry,

$$W = 2\ln\big((1+\gamma)y_1\big) + 2(1+\gamma)(M - y_1).$$

FOC in y_1: $\dfrac{2}{y_1} = 2(1+\gamma)$, i.e., $y_1^o = \dfrac{1}{1+\gamma}$.

Step 4. Conclusion. The Nash equilibrium ($y_1 = 1/(1+\gamma)$) coincides with the social optimum for *every* value of γ. Efficient!

Answer. *Part (e): why?* In (b) the private good ($\gamma_2 = 0$) is treated as a full-cost expenditure by each consumer, but the public good ($\gamma_1 = \gamma > 0$) generates uncompensated benefits for the other consumer — an asymmetric externality that drives the free-riding wedge. In (d) both goods share the *same* externality parameter, so the "cost" of one dollar spent on good 1 and the "opportunity cost" of the alternative dollar spent on good 2 are symmetric across the two consumers. The consumption externalities are balanced — what one consumer confers on the other via good 1 she also receives via good 2 — and the private and social optima coincide.

2. **MWG 11.D.4 – Heterogeneous nondepletable externalities.** Consider the nondepletable externality of Section 11.D, with the twist that the externalities produced by the J firms are *not* homogeneous: consumer i's derived utility, given the firms' externality levels $h_1, \dots, h_J$, is $\phi^i(h_1, \dots, h_J) + w^i$. Compare the competitive equilibrium and the efficient levels of $(h_1, \dots, h_J)$. What tax/subsidy scheme restores efficiency? Under what condition should each firm face the same tax rate?

Answer. *Setup.* Each firm j chooses $h_j \geq 0$; its own profit is $\pi_j(h_j)$, increasing in h_j over some range. The externality is *nondepletable*: every consumer i experiences all firms' externality levels simultaneously.

Step 1. Social planner's problem. The planner chooses $(h_1, \dots, h_J)$ to maximize the sum of consumer welfare and firm profits:

$$\max_{h_1,\dots,h_J \geq 0} \sum_{i=1}^{I} \phi^i(h_1, \dots, h_J) \; + \; \sum_{j=1}^{J} \pi_j(h_j).$$

Step 2. Planner's FOC. For each j, differentiating in h_j,

$$\pi_j'(h_j^o) \; + \; \sum_{i=1}^{I} \frac{\partial \phi^i(h^o)}{\partial h_j} \; \leq \; 0, \qquad \text{with equality if } h_j^o > 0.$$

Equivalently, the marginal social value of h_j is zero: firm j's marginal profit plus the sum of the consumers' marginal utilities from h_j (which is *negative* for a negative externality) equals zero.

Step 3. Competitive equilibrium. Each firm j ignores the externality and maximizes its own profit:

$$\pi_j'(h_j^*) \; \leq \; 0, \qquad \text{with equality if } h_j^* > 0.$$

The FOC contains no ϕ^i term because the firm doesn't internalize consumer effects.

Step 4. Wedge and the corrective tax. Compare the two FOCs at the socially optimal h^o. To make firm j's private FOC coincide with the social FOC, the firm's private objective must be reduced by an amount whose marginal value in h_j equals $\sum_i \partial\phi^i/\partial h_j$. A per-unit tax

$$\boxed{t_j = -\sum_{i=1}^{I} \frac{\partial \phi^i(h^o)}{\partial h_j}}$$

does exactly that. Because ϕ^i is decreasing in h_j (for a negative externality), the sum inside is negative, so $t_j > 0$: a tax.

Step 5. Common tax rate? All firms should face the same rate iff

$$\sum_{i=1}^{I} \frac{\partial \phi^i(h^o)}{\partial h_j} = \sum_{i=1}^{I} \frac{\partial \phi^i(h^o)}{\partial h_k} \qquad \text{for every pair } j, k.$$

That is, only if the marginal social damage of each firm's externality is the same. This holds automatically if firms' externalities enter ϕ^i symmetrically (a common "pollution index"), but not if consumers care differently about, say, particulate matter vs. noise emissions.

Intuitively, the Pigouvian principle applies firm by firm: tax each firm at exactly the marginal social damage *it* imposes. A uniform tax rate is efficient only when all firms happen to impose the same marginal damage — a knife-edge case that doesn't generalize once externalities are heterogeneous.

3. **MWG 11.D.7 – Neighborhood choice with prestige externalities.** A continuum of individuals choose between neighborhoods A and B. Building a house costs c_A in A and $c_B < c_A$ in B. Each individual i has a prestige parameter $\theta_i \in [0, 1]$ (uniformly distributed across the population). The prestige of neighborhood k is $\bar{\theta}_k$, the average θ of its residents. An individual with prestige θ living in neighborhood k has net utility

$$u(\theta, k) = (1 + \theta)(1 + \bar{\theta}_k) - c_k.$$

More-prestigious individuals value a prestigious neighborhood more (a Spence-style single-crossing structure). Assume $c_A, c_B < 1$ and $c_A - c_B \in (\frac{1}{2}, 1)$.

   (a) Show that in any equilibrium both neighborhoods must be occupied.

   (b) Show that in any equilibrium with $\bar{\theta}_A \neq \bar{\theta}_B$, there is a cutoff $\hat{\theta}$ such that types $\theta \geq \hat{\theta}$ live in A and types $\theta < \hat{\theta}$ live in B. Characterize $\hat{\theta}$.

   (c) Show that a small change in $\hat{\theta}$ (with transfers) produces a Pareto improvement.

Answer. *Part (a): both neighborhoods must be occupied.*

Suppose for contradiction that only one neighborhood is occupied. Consider the most prestigious individual ($\theta = 1$).

Step 1. If only B is occupied ($\bar{\theta}_B = 1/2$). Her utility is $(1+1)(1+\frac{1}{2})-c_B = 3-c_B \le 3$. If she deviated to A (which she would find empty, so $\bar{\theta}_A = 1$ upon her arrival — treating herself as the sole resident), her utility would be $(1+1)(1+1)-c_A = 4-c_A > 3 \ge 3-c_B$. Deviation is strictly profitable, contradicting equilibrium.

Step 2. If only A is occupied ($\bar{\theta}_A = 1/2$). Analogous argument: her utility in A is $3 - c_A$, in an empty B (with $\bar{\theta}_B = 1$) it is $4 - c_B > 3 - c_A$. Again a strict deviation.

Hence in equilibrium both neighborhoods have positive population.

Answer. *Part (b): sorting equilibrium with a prestige cutoff.*

Step 1. Single-crossing. Fix any two types θ, θ' with $\theta > \theta'$. Their utility differences between neighborhoods are

$$\Delta u(\theta) \equiv u(\theta, A) - u(\theta, B) = (1+\theta)(\bar{\theta}_A - \bar{\theta}_B) - (c_A - c_B),$$
$$\Delta u(\theta') = (1+\theta')(\bar{\theta}_A - \bar{\theta}_B) - (c_A - c_B).$$

If $\bar{\theta}_A > \bar{\theta}_B$, then $\Delta u(\theta) - \Delta u(\theta') = (\theta - \theta')(\bar{\theta}_A - \bar{\theta}_B) > 0$: the higher-prestige type has a strictly greater incentive to live in A.

Step 2. Cutoff structure. By the single-crossing property, the set of types who prefer A is upward-closed: if type θ' prefers A, so does every $\theta > \theta'$. Hence there is a cutoff $\hat{\theta}$ such that types $\theta \ge \hat{\theta}$ live in A and types $\theta < \hat{\theta}$ live in B.

Step 3. Compute the cutoff. Given the cutoff, the averages are $\bar{\theta}_A = \frac{1+\hat{\theta}}{2}$ and $\bar{\theta}_B = \frac{\hat{\theta}}{2}$ (each is the midpoint of the sub-uniform on $[\hat{\theta}, 1]$ and $[0, \hat{\theta}]$, respectively). The marginal type $\hat{\theta}$ is indifferent between the two neighborhoods:

$$(1+\hat{\theta})\left(1 + \frac{1+\hat{\theta}}{2}\right) - c_A = (1+\hat{\theta})\left(1 + \frac{\hat{\theta}}{2}\right) - c_B.$$

Simplifying,

$$(1+\hat{\theta}) \cdot \frac{1}{2} = c_A - c_B \iff \hat{\theta} = 2(c_A - c_B) - 1.$$

Under the assumption $c_A - c_B \in (\frac{1}{2}, 1)$, $\hat{\theta} \in (0, 1)$, so the cutoff lies strictly inside the type space and both neighborhoods have positive mass.

Answer. *Part (c): Pareto-improving reallocation.*

The key insight: the sorting equilibrium ignores the *externality* each mover imposes on both neighborhoods (raising or lowering their average prestige). A planner can exploit this.

Step 1. Move a thin slice from A to B. Take a small $\varepsilon > 0$ and move all types in $[\hat{\theta}, \hat{\theta} + \varepsilon]$ from A to B. Both averages *rise*:

- $\bar{\theta}_A$ rises because the moved types were the least prestigious of A's residents;
- $\bar{\theta}_B$ rises because the moved types are the most prestigious of B's residents.

Concretely, both averages rise by $\varepsilon/2$ to leading order.

Step 2. Benefit for non-movers. Every non-mover of type θ enjoys a utility gain of $(1+\theta)\cdot\varepsilon/2$ from the higher average in her neighborhood. Aggregating,

$$B(\varepsilon) = \int_0^{\hat{\theta}} (1+\theta)\frac{\varepsilon}{2}\,d\theta + \int_{\hat{\theta}+\varepsilon}^{1} (1+\theta)\frac{\varepsilon}{2}\,d\theta.$$

Step 3. Cost for movers. A mover of type $\theta \in [\hat{\theta}, \hat{\theta}+\varepsilon]$ was indifferent at $\theta = \hat{\theta}$; her utility loss from moving to B (before the average changes) is $(1+\theta)\cdot\frac{1}{2} - (c_A - c_B)$. Adjust for the newly higher averages: her net utility change is $(1+\theta)\big[-\frac{1}{2}+\frac{\varepsilon}{2}\big] + (c_A - c_B)$. The total cost is

$$C(\varepsilon) = \int_{\hat{\theta}}^{\hat{\theta}+\varepsilon} \Big[(1+\theta)\cdot\frac{1-\varepsilon}{2} - (c_A - c_B)\Big]\,d\theta.$$

Step 4. Compare marginal effects at $\varepsilon = 0$. Differentiating both integrals w.r.t. ε and evaluating at $\varepsilon = 0$,

$$\begin{aligned}
\left.\frac{dB}{d\varepsilon}\right|_{\varepsilon=0} &= \int_0^1 (1+\theta)\cdot\frac{1}{2}\,d\theta = \frac{1}{2}\cdot\big[\theta + \tfrac{\theta^2}{2}\big]_0^1 = \frac{3}{4},\\
\left.\frac{dC}{d\varepsilon}\right|_{\varepsilon=0} &= (1+\hat{\theta})\cdot\frac{1}{2} - (c_A - c_B) = 0,
\end{aligned}$$

where the second equality uses the cutoff-defining condition from part (b).

Step 5. Conclude. $dB/d\varepsilon > dC/d\varepsilon$ at $\varepsilon = 0$, so a small move produces strictly positive net benefit. With transfers from non-movers (who gain from the higher averages) to movers (who lose from crossing to the more-expensive-per-utility neighborhood, though the loss is second-order), a Pareto improvement is achievable.

Intuitively, the equilibrium sorting is inefficient because each individual chooses a neighborhood based only on *her own* indifference condition, ignoring the prestige externality her presence has on *everyone else*. The marginal mover values the neighborhoods equally (that's the cutoff condition), so shifting her costs her nothing at the margin — but her move raises the average prestige of both neighborhoods, which *everyone else* values. The planner picks up a first-order gain in aggregate welfare at a second-order cost to the movers, and transfers redistribute so all types are strictly better off.

4. **Complement (aggregating individual preferences).** Exercises 1–3 above studied *aggregation externalities*: individuals' choices affect a public good (contribution levels, pollution, neighborhood prestige) that everyone else consumes. This complement pulls

the aggregation question in a different direction — *can two individuals' preferences even be aggregated coherently into a third person's preference?*

Renata and Sergio each have complete, transitive preferences $\succsim_R$ and $\succsim_S$ over a finite set X. Talia has "preferences for disagreement": she likes to end up on the opposite side of Renata and Sergio when they agree, and is indifferent when they disagree. Formally, for any $x, y \in X$,

- if $x \succsim_R y$ *and* $x \succsim_S y$, then $\neg(x \succsim_T y)$;
- if $\neg(x \succsim_R y)$ *and* $\neg(x \succsim_S y)$, then $x \succsim_T y$;
- if exactly one of $x \succsim_R y$, $x \succsim_S y$ holds, then $x \sim_T y$.

Is $\succsim_T$ well-defined? Complete? Transitive?

Answer. *Not well-defined in general.*

Step 1. Where the trouble lies. The mixed clause declares $x \sim_T y$, which requires both $x \succsim_T y$ and $y \succsim_T x$. But the value of $y \succsim_T x$ is fixed *separately* by the rule applied to the pair (y, x), and the two rulings can clash.

Step 2. Concrete conflict. Take $X = \{x, y\}$ with $x \sim_R y$ (Renata is indifferent) and $y \succ_S x$ (Sergio strictly prefers y). Consider both directions:

- Pair (x, y): $x \succsim_R y$ holds (indifference gives both directions) but $x \succsim_S y$ fails. This is the mixed case, so the rule declares $x \sim_T y$, i.e., $y \succsim_T x$ must hold.
- Pair (y, x): $y \succsim_R x$ (indifference again) *and* $y \succsim_S x$ (strict preference includes weak). Both agree, so the rule declares $\neg(y \succsim_T x)$.

The two rulings on the single relation $y \succsim_T x$ clash: the first says it holds, the second says it doesn't. So $\succsim_T$ is not well-defined whenever one of the two agents is indifferent and the other has a strict ranking on the same pair.

Answer. *Even when well-defined, not complete.*

Suppose $\succsim_R$ and $\succsim_S$ are strict linear orders (no indifferences). Then the mixed case never triggers and $\succsim_T$ is well-defined. Check reflexivity: at (x, x), both Renata and Sergio agree ($x \succsim_R x$ and $x \succsim_S x$), so the "agree" clause declares $\neg(x \succsim_T x)$. But reflexivity requires $x \succsim_T x$ — so $\succsim_T$ is not reflexive. A non-reflexive relation cannot be complete on any set that contains at least one element (completeness $\Rightarrow$ reflexivity via $x \succsim x$ or $x \succsim x$).

Answer. *Even when well-defined and reflexive by convention, not transitive.*

With Renata and Sergio's strict orders, the rule reads: if both agree $x \succ y$, Talia *prefers the loser*, $y \succ_T x$; if they disagree, Talia is indifferent.

Step 1. *Construct a cycle.* Take $X = \{x, y, z\}$ with

$$x \succ_R y \succ_R z \qquad \text{and} \qquad y \succ_S z \succ_S x.$$

Step 2. *Pairwise Talia rankings.* Comparing each pair,

- (x, y): Renata says $x \succ y$, Sergio says $y \succ x$: they *disagree*, so $x \sim_T y$.
- (x, z): Renata says $x \succ z$, Sergio says $z \succ x$: they *disagree*, so $x \sim_T z$.
- (y, z): Renata says $y \succ z$, Sergio says $y \succ z$: they *agree*, so the "agree" clause gives $z \succ_T y$.

Step 3. *Transitivity fails.* $x \sim_T z \succ_T y$ but $x \sim_T y$ (not $x \succ_T y$). Transitivity would require the strict comparison $z \succ_T y$ combined with the tie $x \sim_T z$ to yield $x \succ_T y$, but the rule delivers only $x \sim_T y$.

Intuitively, "always disagree with the consensus" is a self-undermining aggregation rule. Reflexivity requires an agent to agree with herself, but the disagreement rule forbids exactly that when the two "advisors" agree — including a bundle with itself. Even setting reflexivity aside, chaining a strict Talia preference through a tie breaks transitivity in the standard way. The moral: aggregating individual preferences into a coherent collective preference is much subtler than it looks, and simple "mixing" rules routinely fail — exactly the sort of aggregation pathology that public-goods provision and mechanism design have to grapple with.