Negotiated collective inter-temporal decisions with heterogeneous inter-temporal preferences

Ram Fishman

ROUGH DRAFT. EXTREMELY PRELIMINARY.

Abstract

In making inter-temporal decisions involving public goods (or bads), agents can often all benefit from coordinating a Paretto efficient policy rather than following the non-cooperative Nash equilibrium ("Tragedy of the Commons"). There are numerous Paretto efficient policies, and in negotiations over the choice of one, it is intuitively appealing that impatient agents will have stronger bargaining power, since the benefits of cooperation are accrued in the future, and hence more highly valued by patient agents. This could suggest that the negotiated outcome might be weighted in their favor, and in the absence of side payments, might lead to higher collective impatience. I attempt to formulate this notion in the framework of cooperative models of bargaining, and show it holds in a simple model of global warming.

Introduction

The analysis of commons problems typically compares socially optimal consumption paths, which are Pareto optimal solutions, and Markov Perfect Equilibria (MPEs), which represent the non-cooperative solutions. The tragedy of the commons in the dynamic setting is the sub-optimal, and, in some sense (but not always), the less patient nature of the non-cooperative solution. Lack of patience is reflected in the lower collective investment in the future in the form of common capital, resource stocks or avoidance of public bads.

Socially, it is of course desirable to realize one of the Paretto optimal solutions. An important question is what mechanisms are able to achieve this goal? For example, can reversion to some MPE as a form of punishment be able to sustain the social optimum as a SPE within the dynamic game? Or perhaps there are political mechanisms external to the game which can enforce some types of socially beneficial strategies (i.e. linkage)? Here, I will assume that such linkage exists, so that whatever policy is agreed upon is enforceable, and ask instead which policy is actually going to be chosen. There is, of course, no simple answer to this question. Nash, in his classical paper, Nash (1950), described the bargaining problem as follows:

"A two-person bargaining situation involves two individuals who have the Opportunity to collaborate for mutual benefit in more than one way "(p. 155).

The bargaining problem has never received a simple, universal, agreed, clear solution in economic theory. Nash's 1950 analysis defined an approach, called 'cooperative' which abstracted the problem from the details of the bargaining process itself. Other approaches aimed to base the solution on non-cooperative game theory and modeled the details of the bargaining process. Various solutions were thus arrived at over the years. Although they often disagree on both the formulation and the exact outcome, one component that often plays a central role is the disagreement point/outcome – a benchmark of utilities that represent the event of failure to reach any agreement. Most of the cooperative formulations of the problem also try to capture the idea of bargaining strength in their formulation of the problem.

In what follows, I try to investigate the following intuitive idea: if agents differ in the degree to which they value the future, they will also differ in the degree to which they benefit from cooperation, since the benefits of the latter, in contrast to the non-cooperative behavior, are accrued in the future. If that is the case, we may expect the more patient agents to have less of a bargaining power in the negotiations, so the negotiated outcome may be titled towards the impatient agents.

I choose to abstract away from any particular mechanism, and pursue the idea with cooperative modeling, because this expectation is not based on any particular mechanism of negotiations, but seems quite general and reasonable whenever bargaining strength is allowed to play a role. I only assume that an efficient mechanism is at play, which ensures that an efficient policy is adopted. Clearly, every other policy is Pareto-dominated by an efficient one, so rational agents will choose one of these efficient policies if the use an efficient negotiations mechanism. These various Paretto-efficient policies can be characterized as those which optimize different weighted averages of the agents' payoffs. Thus, a negotiated outcome tilted towards impatient agents will be one that maximizes an average of payoffs with weights concentrated with impatient agents. As such, it is also likely to be a collectively more impatient policy: in the case of extraction of an exhaustible resource, for example, this could mean that the aggregate extraction rate is higher than that of the evenly weighted optimal policy. However, if side payments are allowed as a form of compensation, the result of the bargaining will likely revert to that of the maximization of the un-weighted mean of payoffs. The implication is, then, that the allowance of side transfers may result in a more patient collective behavior. In the context of climate change, for example, allowing side payment as part of the post-Kyoto negotiations may lead to lower aggregate global emissions of GHGs.

Why should we expect to find heterogeneity in agents' time preferences? Both theoretically and empirically, there seem to be ample reason to believe that time preferences can be highly heterogeneous across agents. The empirical investigation of pure time discounting rates contains wide spectrums of estimates which differ across individuals, societies, time horizons and situations. In particular, poorer individuals and societies tend to discount the future less than rich ones.
Also, in commons problems, it may well be that agents have alternative investment strategies (alternative to investment in the common, that is) which differ in their rate of returns. One possible example is climate change – rates on returns on investments may be much higher in certain developing countries in which economic growth is accelerating. Also, one could argue that better governance is manifested, among other things, in greater foresight, and that some poorly governed states therefore tend to more myopic.

Hypothesis Formulation

As described above, there is no canonical single description of the bargaining problem. In making a general argument, however, it seems to make more sense to follow the general treatment of cooperative models of the bargaining problem, since it is impossible to model in a general way the precise bargaining process, which is often a complex political process.[1] But even when abstracting away from the details of the process, there are still numerous cooperative proposed solutions to the problem. In which of these should we try to test the naïve argument presented above? Let us focus on the case of two agents. Most formulations of the bargaining problem include a set S of feasible alternatives, preferences for both agents over this set, and the disagreement point d. Among the many candidate solutions, the review by Thomson (1994) identifies three formulations as being particularly robust and justifiable: “Nash's original solution, which is axiomatically derived, and selects the point of S at which the product of utility gains from d is maximal, a solution due to Kalai and Smorodinsky (1975), which selects the point of S at which utility gains from d are proportional to their maximal possible values within the set of feasible points dominating d, and the solution that simply equates utility gains from d, the Egalitarian solution.” I will focus on the egalitarian and the Nash solution in this draft.

Another important caveat to keep in mind is that there is, in general, no single MPE that can define the disagreement point. Moreover, not all MPEs are in fact, more impatient (Dutta & Sundaram 93), so we can’t expect to formulate a general statement without some assumptions on the MPE chosen.

An important difference between the Nash and the Egalitarian solution is that the latter makes a comparison between the payoffs of two agents. In this context, this would entail comparing the inter-temporal payoffs of two agents with different time discount rates. Let these payoffs be given as

(A) $\mathrm{V}_\mathrm{i} = \sum_{t=0}^{\infty} \delta_\mathrm{i}^\mathrm{t} v_i(t)$,

where $v_i(t)$ is the utility of agent i at time t. Arguably, when comparing two such inter-temporal payoffs, it is more appropriate to use the mean

(B) $\bar{\mathrm{V}}_\mathrm{i} = (1 - \delta_\mathrm{i}) \sum_{t=0}^{\infty} \delta_\mathrm{i}^\mathrm{t} v_i(t)$.

For example, if both agents have a constant in time utility stream v, using (B) means their inter-temporal payoffs are also equal to one another, and using (A) means that even when

[1] In some of the non-cooperative models of bargaining, time discounting plays a role in the bargaining process itself (perhaps most notable is Rubinstein 1987). The time discounting which plays a role in the tragedy of the commons, however, is over time scales that might be much longer than those of the bargaining process itself. It seems to make better sense not to mix these two notions of time discounting.

both agents enjoy the same constant utility at each period, the overall payoff of the more patient agent is higher. This seems problematic. On the other hand, if the discount rate is interpreted as an interest rate, it might be reasonable to use (A) definition.

The distinction is immaterial to the Nash solution, but in the egalitarian solution can be important. Compared to (B), the definition (A) is biased towards the patient agent. If we prove that the egalitarian solution is tilted towards the impatient agent when using (B) to define inter-temporal payoffs, it will therefore also hold for the definition (A) as well. This is what I find in the single calculated example in this draft, in the last section.

Generally, suppose that non-cooperative behavior is represented by an MPE whose payoff to player i is V_i^* (and V is defined by (A) or (B) above). Denote the payoff to player i when the Paretto-efficient strategy maximizing $xV_1 + (1-x)V_2$ is $V_i(x)$. Clearly, $V_1(x)$ is increasing in x, and $V_2(x)$ decreasing in x.

We expect collective behavior to become more impatient as the weight assigned to the impatient agent increases. Impatience can be measured in aggregate extraction, emissions, etc… depending on the specific problem. This is generally not hard to prove in any specific example.

Let δ_i be the agents' discount rates. Assume agent 1 is more patient, so that $\delta_1 > \delta_2$. I formulate the ideas presented in the introduction as a hypothesis that negotiations will lead to a choice of the strategy which maximizes $xV_1 + (1-x)V_2$ with x<1/2.

Let $\Delta_k(x) = V_k(x) - V_k^*$. If we adopt the egalitarian solution, this hypothesis translates to

(HEa) $\Delta_1(x) = \Delta_2(x)$ at x<1/2 Using definition A
(HEb) $\Delta_1(x) = \Delta_2(x)$ at x<1/2 Using definition B

As discussed above, if (HEb) holds, (HEa) does as well. If we adopt the Nash solution, the hypothesis translates to:

(HN) $\text{Argmax}\{\Delta_1(x) \times \Delta_2(x)\} < 1/2$

Alternatively, we might also expect that the set of Paretto optimal strategies which is favorable for both players will be asymmetrically supported, i.e. that

Inf S < 1-sup S, where $S = \{x \mid \Delta_1(x) > 0, \Delta_2(x) > 0\}$

Testing the conjecture for a Simple Dynamic Model of global Warming

It is rarely possible to obtain closed form solutions for the present discounted inter-temporal utility streams of dynamic commons games. Even when such closed form solutions exist for both the MPE and the social optimum under a common time discounting, it is hard to obtain them for asymmetric time discounting, because the collective problem is no longer captured by a dynamic programming Bellman equation. Thus, it is hard to obtain a sample of closed form solutions with which to test this idea.

Dutta and Radner (2004) offer a commons dynamic game model of global warming, which allows for relatively simple closed form expressions for the quantities appearing in the above conjectures. I use a simplified setting of their model, in which there are two countries which are identical except for the time discounting rates. Let GHG emissions be $a_i(t)$ i=1,2, $A(t) = a_1(t) + a_2(t)$ be total GHG emissions, and suppose that the equation of motion for g(t), the amount of GHG in the atmosphere, is $g(t+1) = \sigma g(t) + A(t)$. The utility of country i at time t is given by

(1) $$v_i(t) = h[a_i(t)] - cg(t)$$

where h(a) is concave production function of emission levels, and damages are assumed to be linear. Total payoff is given by

(2a) $$\mathrm{V_i} = \sum_{t=0}^{\infty} \delta_i^t v_i(t)$$

or

(2b) $$\mathrm{V_i} = (1-\delta_i) \sum_{t=0}^{\infty} \delta_i^t v_i(t).$$

I will use the form (2b), since, as discussed above, all results apply to the form (2a) if they hold for (2b).

Given constant emission levels $a_i(t) = b_i, b = b_1 + b_2$, and initial GHG levels g, payoffs are easily calculated to be:

(3) $$\mathrm{V_i}(b_1, b_2, \mathrm{g}) = \mathrm{h}(b_i) - b\,\frac{\delta_i c}{1-\delta_i\sigma} - g\,\frac{(1-\delta_i)\, c}{1-\delta_i\sigma}$$

The first term represents the average per period benefit of emissions, the second represent the average damages of per period emissions, projected over the remainder of time, and the third represents the damages from the initial concentration of GHGs. For details, see Dutta and Radner.

Both the Paretto efficient (GPO) policies and the Markov perfect Equilibrium (MPE) policy which I use here to describe the non-cooperative behavior consist of constant emission levels. MPE emissions a_i^* satisfy:

(4a) $h'(a_i^*) = \frac{\delta_i c}{1-\delta_i\sigma}$

And GPO emissions, $\widehat{a_k}$, which maximize the weighted aggregate utility $x_1V_1 + x_2V_2$, satisfy:

(4b) $x_k h'(\widehat{a_k}) = x_1 \frac{\delta_1 c}{1-\delta_1\sigma} + x_2 \frac{\delta_2 c}{1-\delta_2\sigma}$

(I neglect to include the weights in the notation $\widehat{a_k}$ but this should be clear from the context). Notice that $\widehat{a_k} < a_k^*$ because of the concavity of h(a). Also, if $\delta_1 > \delta_2$, $a_1^* < a_2^*$ (patient country pollute less).

Denote $\beta_i = \frac{\delta_i c}{1-\delta_i\sigma}$. Each country's gain from cooperation (with weights x_i), is therefore:

(5) $\Delta_k(x_1,x_2) = V_k(\widehat{a_1},\widehat{a_2},g) - V_k(a_1^*,a_2^*,g) = (h(\widehat{a_k}) - \hat{a}\beta_i) - (h(a_k^*) - a^*\beta_i)$

Notice that β_i is increasing in δ_i, and maintains the ordering of impatience among the countries. Let us assume that country 1 is the more patient country, so that $\beta_1 > \beta_2$. I now show that both the egalitarian and the Nash solutions will predict a negotiated outcome with $x_1 < x_2$, i.e. a policy which is tilted towards the impatient country, by showing that:

Proposition: (a) $\Delta_1(\frac{1}{2},\frac{1}{2}) > \Delta_2(\frac{1}{2},\frac{1}{2})$

(b) $\frac{\delta[\Delta_1(\frac{1}{2}+\varepsilon,\frac{1}{2}-\varepsilon)\Delta_2(\frac{1}{2}+\varepsilon,\frac{1}{2}-\varepsilon)]}{\delta\varepsilon} < 0$

Proof:

(a) If $x_1 = x_2$, then $\widehat{a_1} = \widehat{a_2}$, so that $\Delta_1 - \Delta_2 = (a^* - \hat{a})(\beta_1 - \beta_2) - [h(a_1^*) - h(a_2^*)]$. Now, $a^* > \hat{a}$ (tragedy of the commons), and $a_1^* < a_2^*$, so we easily conclude that $\Delta_1 > \Delta_2$.

(b) When the weights change, only the first term in (5) changes. Using (4b), we see that $d\Delta_1 = h'(\widehat{a_1})d\widehat{a_1} - \beta_1 d\widehat{a_1} - \beta_1 d\widehat{a_2} = \frac{x_2}{x_1}\beta_2 d\widehat{a_1} - \beta_1 d\widehat{a_2}$, and similarly $d\Delta_2 = \frac{x_1}{x_2}\beta_1 d\widehat{a_2} - \beta_2 d\widehat{a_1}$, so that when $x_1 = x_2$, $d\Delta_1 = -d\Delta_2$. It follows that $d[\Delta_1\Delta_2] = d\Delta_1 \times \Delta_2 + d\Delta_2 \times \Delta_1 = (\Delta_1 - \Delta_2)\, d\Delta_2$. Clearly, if x_1 increases and x_2 decreases, the utility of country 2 in the cooperative policy decreases, so Δ_2 decreases as well, and hence $d[\Delta_1\Delta_2] < 0$, because as shown in (a), $\Delta_1 - \Delta_2 > 0$.

Part (a) of the proposition proves the egalitarian solution occurs at a weight $x_2 > \frac{1}{2}$, whereas part (b) shows that in order to increase $\Delta_1\Delta_2$, x_2 should be moved upwards from the equal distribution, so that is where the Nash solution will be placed.

In closing, it is worth noting that Dutta and Radner consider a different form of asymmetry in which the cost c is different for each agent. They show that as the cost of one agent increases (and the other's decreases), the span of incentive compatible strategies of that agent decreases (and the other's increases; compared to the reversion to the MPE). Looking at the expressions above, it is apparent that increasing c is identical to an increase in the discount factor (except for the factor in 2b). So we may interpret their result, in the context of this draft, as suggesting an increase in the set of options attractive to an agent whose impatience is decreasing. This is in the spirit of the notion presented at the end of section 2, and in the general spirit of the ideas discussed here.

References

Ken Binmore, Martin J. Osborne and Ariel Rubinstein: Noncooperative models of bargaining. Handbook of Game Theory with Economic Applications, Volume 1, 1992, Pages 179-225

William Thomson : Cooperative models of bargaining. Handbook of Game Theory with Economic Applications, Volume 2, 1994, Pages 1237-1284

Dutta, Prajit K & Sundaram, Rangarajan K, 1993. "The Tragedy of the Commons?" Economic Theory, Springer, vol. 3(3), pages 413-26, July.

Nash, John, 1950."The Bargaining Problem," Econometrica, Econometric Society, vol. 18(2), pages 155-162, April.

Rubinstein, Ariel, 1982."Perfect Equilibrium in a Bargaining Model," Econometrica, Econometric Society, vol. 50(1), pages 97-109, January.

PK Dutta, R Radner: “A strategic analysis of global warming: Theory and some numbers”.