nature climate change

SUPPLEMENTARY INFORMATION
DOI: 10.1038/NCLIMATE2229

ADDENDUM

A bargaining game analysis of international climate negotiations

Rory Smead, Ronald L. Sandler, Patrick Forber and John Basl

Nature Climate Change **4,** 442–445 (2014); published online 11 May 2014; addendum published after print 8 August 2014.

In the Supplementary Information file for this Letter, the authors have added a section titled "Description of simulation procedure", which provides a detailed account of the primary simulation procedure used in the study. This section is intended to allow others to directly replicate the results, and contains some detail and accompanying explanation omitted from the original version. The revised version also includes notational corrections to the proof of Proposition 1.

A Bargaining Game Analysis of International Climate Negotiations

Rory Smead, Northeastern University
Ronald Sandler, Northeastern University
Patrick Forber, Tufts University
John Basl, Northeastern University

August 4, 2014

Supplemental Material

The model and definitions

Here we add details about the model, compare our model to other standard approaches to learning and bargaining, and provide supplemental results.

Negotiations are represented as iterative interactions in a bargaining game with N players. Between each set of demands, players can change their demand in an attempt to maximize their payoff relative to the mean of past demands from other players. This learning rule is a variation on the simultaneous fictitious play learning rule between interactions (Nash, 1950; Brown, 1951; Fudenberg and Levine, 1998; Young, 2004).

Each player has a continuous strategy set $strat_i = \{d|d \in [\delta Max_i, Max_i]\}$ this represents the set of possible demands in negotiation. The default value for Max_i is 1. Where an agents maximum strategy is restricted by a factor r, $Max_i = 1 - r$ and when the restriction is also applied to the minimum demands, $(\delta Max_i + r)$ is the lower-bound on demands. Each player has an associated size s_i. Let $t = \sum_i d_i s_i$ be the total demands from all players and let $T \in [0, 1]$ represent the target value for negotiation. The payoff function for each player is given by:

$$\pi_i(d_i, t) = \begin{cases} d_i & \text{if } t \leq T \sum_i s_i \\ \delta d_i & \text{if } t > T \sum_i s_i \end{cases} \tag{1}$$

A best response is defined as usual:

$$Br_i(t) = \{d_x | \pi_i(d_i, t) \geq \pi_i(d'_i, t) \text{ for all } d'_i \in strat_i\}. \tag{2}$$

And a Nash equilibrium is a strategy profile $(d_1, ..., d_n)$ such that $d_i \in Br_i(t)$ for all d_i. A Nash equilibrium is strict if any other strategy by any player delivers a strictly worse payoff to that player.

A successful negotiation will be one that achieves a feasible solution defined as a set of demands where all players are better off than their max demand out of agreement. This amounts to:

Definition 1. *A set of demands $(d_1, ..., d_n)$ is a feasible solution if and only if it $T\sum_i s_i \geq \sum_i d_i s_i$ and $d_i > \delta Max_i$ for all i.*

Additionally, a disagreement outcome will be a set of demands such that agreement cannot be facilitated by any single player changing her strategy. More precisely:

Definition 2. *A set of demands $(d_1, ..., d_n)$ is a disagreement outcome iff $T\sum_i s_i < (\sum_i s_i d_i) - (s_j d_j - s_j \delta Max_j)$ for all players j.*

Strict Nash equilibria

It is possible to prove some general results about the equilibria of the game. More specifically, there are both feasible solutions and disagreement equilibria that are strict Nash equilibria. Strict Nash equilibria form strongly stable states for a wide range of learning dynamics (Weibull, 1995).

Proposition 1. *Feasible solutions are strict Nash equilibria iff $t = T\sum_i s_i$.*

Proof. Proof proceeds by cases.

(i) Suppose $(d_1, ..., d_n)$ is a feasible solution such that $t = T\sum_i s_i$. Without loss of generality, let d_j represent the demand of player j. Now consider an alternative strategy for j, d'_j. If $d'_j < d_j$, then the new total is $t' = t-(d_j - d'_j)s_j$ and the payoff to j is now $\pi_j(d'_j, t') = d'_j < d_i = \pi_j(d_j, t)$. If $d'_j > d_j$ then $t' = t + (d'_j - d_j)s_i$ and $\pi_j(d'_j, t')) = \delta d'_j$ and since $(d_1, ..., d_n)$ is a feasible solution $\pi_j(d'_j, t')) < \pi_j(d_j, t)$. Therefore, $(d_1, ..., d_n)$ is a strict Nash equilibrium.

(ii) Suppose that $(d_1, ..., d_n)$ is a strict Nash equilibrium and is a feasible solution ($t \leq T\sum_i s_i$). Without loss of generality, let d_j represent the demand of player j. If $t < T\sum_i s_i$, then the strategy $d'_j = d_j + (T\sum_i s_i - t)/s_i$ is such that $\pi_j(d'_j, t') = d'_j > d_j = \pi_j(d_j, t)$. So, $(d_1, ..., d_n)$ would not be a strict Nash, hence $t \geq T\sum_i s_i$. □

Proposition 2. *If $(Max_1, ..., Max_n)$ is a disagreement outcome, then it is a strict Nash equilibrium.*

Proof. Suppose $(Max_1, ..., Max_n)$ is a disagreement outcome, then $T\sum_i s_i < (\sum_i s_i Max_i) - (s_j Max_j - s_j \delta Max_j)$ for all players j. Without loss of generality let d_j represent the demand of player j. Now consider an alternative strategy for j, $d'_j < Max_j$ and an associated alternative total t'. Note $d'_j \geq \delta Max_j$ by definition of the strategy set. So, $t' > T\sum_i s_i$ by supposition and $\pi_j(d'_j, t') = \delta d'_j < \pi_j(d_j, t) = \delta d_j$ and j is strictly worse off with d'_j. Therefore, $(Max_1, ..., Max_n)$ is a strict Nash equilibrium. □

Learning dynamics and simulations

To investigate whether and how frequently the various equilibria are reached by agents, we employ simulations of adaptive agents in negotiations. Agents are given a determinate size s_i at the beginning of the simulations. Negotiations are modeled as a series of rounds where simultaneous demands are made in the bargaining game. Initial states are determined by a random variable drawn from $[\delta Max_i, Max_i]$ with a uniform distribution. After each round, the demands of all agents are available to all others. Where the strategies are restricted, the initial states are drawn from a uniform distribution in the adjusted strategy set.

Beginning in the second round, agents form expectations of the other agents behavior by examining their history of play. The estimate for an agent's future is the mean of all past demands from that agent. Using these estimates each agent determines what their best-response demand is for the next round. In the next round, they make that demand, observe the demands of all other agents and update their expectations.

Simulations (written in C) are run for 100 rounds. If the expected demand of all players yields a total within 1% of the target value or less ($t \leq T\sum_i s_i + .01$), the negotiation is deemed to be successful. The data points presented in the main text are averaged from 10,000 independent simulations.

Description of simulation procedure

Initializing the simulation: Each agent i in a set of players is assigned a positive real number s_i to denote player size. They are then assigned a closed-interval $[\delta Max_i, Max_i] \subseteq [0, 1]$, which represents the player's strategy set. The parameter $\delta \in (0, 1)$ represents the payoff discount in the event of a failed negotiation.

For each independently simulated negotiation players are assigned initial demands randomly. Let d_i^r denote the demand by player i in round r. d_i^0 represents the pre-negotiation state and is determined by a random variable drawn from a uniform distribution over player i's strategy set.

Simulated rounds of negotiation: Negotiations occur in a series of rounds. Each round consists of a set of demands from each player.

When negotiations begin $d_i^1 = d_i^0$. For each $r \geq 2$, d_i^r is i's best response to the expected demands of others in round r. The estimate for an agent's future demand is the mean of past demands (including the pre-negotiation state d_i^0). Let $E(d_i^r)$ be the expected demand of agent i at round r:

$$E(d_i^r) = \sum_{n=0}^{r-1} d_i^n / r. \tag{3}$$

For each $r \geq 2$, each agent i calculates the total expected demands of all other players $j \neq i$ for that round:

$$E(t_i^r) = \sum_{j \neq i} E(d_j^r) s_j. \tag{4}$$

Each agent then compares $E(t_i^r)$ to the target value (T) weighted by player size:

$$T \sum_k s_k - E(t_i^r). \tag{5}$$

This determines what is in effect the remainder of the available goods given the expected demands of other players. The best response for player i is to demand this remainder adjusted by s_i provided that demand falls within i's strategy set:

$$d_i^r = \begin{cases} (1/s_i)(T\sum_k s_k - E(t_i^r)) & \text{if } (1/s_i)(T\sum_k s_k - E(t_i^r)) \in [\delta Max_i, Max_i] \\ Max_i & \text{otherwise.} \end{cases} \tag{6}$$

Note that $d_i^1 = d_i^0$ and both are included when agents are calculating expected demands. This means that initial demands effectively receive additional weight compared with subsequent demands. The reason for this assumption is that in climate negotiations much is known about the initial conditions and attitudes of other agents. Thus, it reasonable to include the pre-negotiation state d_i^0 that leads to initial demands as part of an agents estimate of others. Furthermore, while this assumption does not alter the qualitative results discussed in the article, it does tend to increase the success rates of negotiations, making the effects of various factors considered more apparent.

Stopping the simulation: If, at any point, the actual demand of all players is equivalent to their expected demands, there will be no subsequent change of demands and the simulation is terminated. If this condition is not met, simulations are terminated after 100 rounds of negotiation.

Classifying the result: After the final round of negotiation (r=100), the cumulative expected demands of all players, $\sum_i E(d_i^{101})s_i$, is compared to the global target value, $T\sum_i s_i$. If this expectation is within 0.01 of the global target value, the negotiation is deemed successful. Otherwise, it is deemed a failure.

Traditional fictitious play

The learning rule described above differs from traditional fictitious play, which does not form expectations by taking averages of opponents play, but by assuming opponents are playing mixed strategies. A mixed strategy is one that randomizes among pure strategies and is represented with a probability distribution over pure strategies. Traditional fictitious play assumes that opponents are playing a fixed mixed strategy where the probability of playing a given pure strategy is equal to the frequency of the opponent using that strategy relative to all past play. Agents then choose a best response—a strategy that maximizes expected utility—relative to their opponents' expected mixed strategies.

Traditional fictitious play faces a number of complications within our model. First, there is a continuous strategy space, which means that as the number of rounds of negotiation increases, calculating best responses becomes computationally intractable. Second, agents using fictitious play are known to occasionally fall into alternating patterns of miscoordination. In such patterns, both players attempt to simultaneously accommodate the other player, then switch back to accommodate again, and fall into an indefinite cycle (Shapley, 1964).

Our modified learning rule has the advantage of avoiding these complications. However, it is also possible to modify the setup of the game in order to investigate traditional fictitious play. First, we can discretize the strategy space, meaning players can only make demands at certain intervals (we will use 0.1, 0.2,...,1.0). This makes the computational tasks of the agents much more tractable. Second, we can introduce some "inertia" into the learning process where, with some probability, individuals simply do not change their behavior after a given round of play (we will use a value of 0.2 for this probability in simulations). This randomness will interupt possible inefficient cycling patterns.

Despite the differences between our modified learning rule and traditional fictitious play, the two learning rules generate qualitatively similar results (see the supplemental results below). The traditional rule is slightly less successful and favors smaller players in successful agreements. Additionally, even with the discretized strategy space, the computations in simulations are significantly more complicated than our modified rule, making it more difficult to examine situations with large numbers of players.

Comparison of modified and traditional fictitious play

To examine traditional fictitious play in our bargaining model in a tractable way, let the strategy space for each player be: $strat_i = \{0.1, 0.2, ..., 1.0\}$. In cases where $\delta > 0.1$, demands lower than δ were removed. Even with this simplification, calculating opponent mixed strategies and best responses for each agent takes exponentially more computation for each additional agent in the game. For this reason, we were only able to simulate models of five or fewer agents. Even in this limited setting, however, the results were qualitatively similar to those of the modified learning rule presented in the main text. In all simulations presented here, it was assumed that on 20% of rounds, agents would not update behavior and maintain their previous demand. In cases where there were multiple best-response options, the lower demand was always preferred.

Many of the results for traditional fictitious play qualitatively similar to those of the modified version that are reported in the main text. As the number of players increased, it dramatically reduced the rate of achieving a successful agreement (Table 1). Heterogeneity of players tends to help reaching an agreement, as fewer players makeup larger proportions of potential total demands success becomes easier to achieve (Table 2). Examining a three-player setting with players of various sizes reveals that reduction burdens fall disproportionally on larger players (Table 3). Restricting initial bargaining demands enhances the prospects for reaching successful agreements (Table 4). Finally, traditional fictitious play also yields the result that prior reductions are most effective if done by a larger number of smaller players as opposed to fewer larger players (Table 5).

There are a few differences between the two learning rules that are worth noting. First, the modified version of fictitious play typically outperforms the traditional version in this bargaining game, being more likely to achieve an agreement. Second, with respect to the effect of prior agreements on negotiations, the traditional rule is far less sensitive to potential increases in minimum demands than the modified version (Table 5). Finally, the distribution of burden in reductions tends to fall much more on the larger players when agents use traditional fictitious play than the modified version, in which results are more equitable (Table 3).

Relation to standard bargaining solutions

Bargaining games have a long and rich history in game theory. The standard approach to bargaining games is to characterize certain equilibria as more desirable, fair, or otherwise salient solutions (Nash, 1950; Kalai and Smorodinsky, 1975; Kalai, 1977). For discussion and overview of bargaining problems and solutions see, e.g., Binmore (1998, 2007); Muthoo (1999).

To compare standard bargaining solutions to learning through traditional and our modified version of fictitious play, consider the simple two player case where $s_1 = 2.0$, $s_2 = 1.0$, $\delta = 0.2$ and $T = 0.5$. Figure 1 shows a plot of the feasible solutions and identifies the various standard bargaining solutions.

Neither traditional fictitious play, nor our modified version converge exactly to any of these specific solutions. Each learning rule produces a wide variety of possible outcomes all along the set of feasible solutions which meet the target value. However, if we examine the average distribution between players, some interesting comparisons emerge. The average solution for traditional fictitious play was approximately 0.406 for player 1 and 0.677 for player 2, relatively close to the Nash bargaining solution. For the modified fictitious play, the average solution was approximately 0.455 for player 1 and 0.590 for player 2, very close to the Kalai-Smorodinsky solution.

Table 1: Rates of successful negotiations as a function of δ and the number of players. Table compares traditional fictitious play with the modified version. 1000 simulations per data point.

	Traditional		Modified	
Condition	$\delta = 0.1$	$\delta = 0.2$	$\delta = 0.1$	$\delta = 0.2$
2 Players	.954	.846	.987	.930
3 Players	.827	.454	.878	.604
4 Players	.674	.226	.717	.385
5 Players	.064	.019	.583	.239

Table 2: Rates of success of traditional fictitious play with four players as Player 1 is varied in size. The other three players are held constant at size 1.0 and $\delta = 0.1$.

Player 1 size	1.0	2.0	3.0
Success of Traditional	.674	.699	.837

Table 3: Average demands in successful agreements for a three player setting with player sizes 1.0, 2.0 and 3.0 with both modified and traditional fictitious play. $\delta = 0.1$. Results rounded to nearest .001.

Player Size	1.0	2.0	3.0
Traditional ($N = 870$)	$\approx .828$	$\approx .521$	$\approx .370$
Modified ($N = 989$)	$\approx .620$	$\approx .528$	$\approx .441$

Table 4: Success rates of negotiations with four agents using traditional fictitious play where initial demands are restricted in various ways. A restriction of r on only maximum initial demands requires that initial demands d such that $0.1 \leq d < 1-r$. Restrictions on minimum initial demands requires $0.1 + r \leq d < 1$. Restrictions on both requires $0.1 + r \leq d < 1 - r$. $\delta = 0.1$.

Degree of restriction	$r = 0$	$r = 0.1$	$r = 0.2$	$r = 0.3$
Restriction on maximum	.674	.805	.835	.798
Restriction on minimum	.674	.642	.538	.417
Restriction on both	.674	.786	.868	.939

Table 5: Success rates of four agents using traditional fictitious play with prior agreements in place. Results include reduction agreements both with and without increased minimum demands. Player sizes: 3.0, 1.0, 1.0, and 1.0. $\delta = 0.2$.

Scenario	Success rate	With increased minimum
No prior reductions	.289	.289
20% reduction from largest	.462	.280
20% reduction from three smallest	.708	.759
10% reduction from all	.534	.559

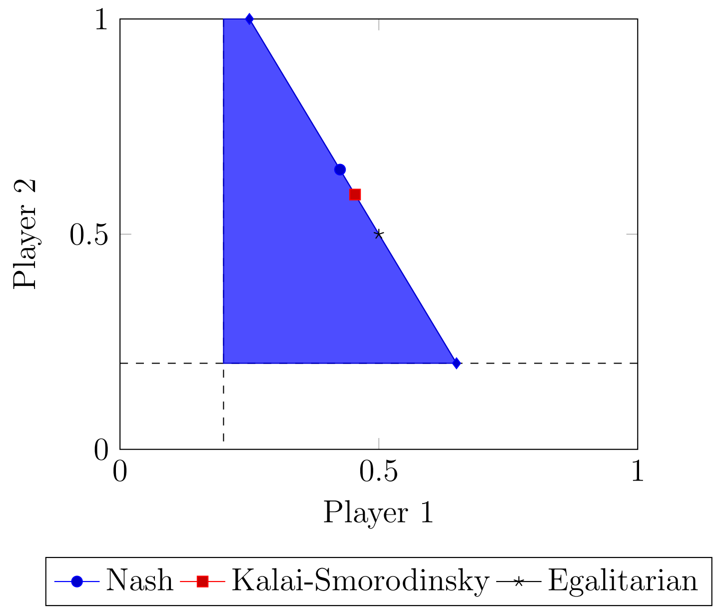

Figure 1: Feasible solution set for two players where $s_1 = 2, s_2 = 1$, $\delta = 0.2$, and $T = 0.5$. The Nash bargaining solution is $(0.425, 0.65)$. The Kalai-Smorodinsky solution is approximately $(0.454, 0.592)$. The egalitarian solution is $(0.5, 0.5)$.

References

Binmore, K. (1998). *Game Theory and the Social Contract, Vol. 2: Just Playing*. Cambridge MA: MIT Press.

Binmore, K. (2007). *Playing for Real*. Oxford: Oxford University Press.

Brown, G. W. (1951). Iterative solutions of games by fictitious play. In T. C. Koopmans (Ed.), *Activity Analysis of Prodcution and Allocation*. New York: Wiley.

Fudenberg, D. and D. Levine (1998). *The Theory of Learning in Games*. Cambridge MA: MIT Press.

Kalai, E. (1977). Proportional solutions to bargaining situations: Interpersonal utility comparisons. *Econometrica 45*, 1623–1630.

Kalai, E. and M. Smorodinsky (1975). Other solutions to nash's bargaing problem. *Econometrica 43*(3), 513–518.

Muthoo, A. (1999). *Bargaining Theory with Applications*. Cambridge: Cambridge University Press.

Nash, J. (1950). The bargaining problem. *Econometrica 18*(2), 155–162.

Shapley, L. S. (1964). Some topics in two-person games. In M. Dresher, L. S. Shapely, and A. W. Tucker (Eds.), *Advances in Game Theory*. Princeton, NJ: Princeton University Press.

Weibull, J. W. (1995). *Evolutionary Game Theory*. Cambridge MA: MIT Press.

Young, P. H. (2004). *Strategic Learning and its Limits*. Oxford: Oxford University Press.