nature climate change

LETTERS

PUBLISHED ONLINE: 2 NOVEMBER 2015 | DOI: 10.1038/NCLIMATE2849

Strategic reasoning and bargaining in catastrophic climate change games

Vilhelm Verendel*, Daniel J. A. Johansson and Kristian Lindgren

Two decades of international negotiations show that agreeing on emission levels for climate change mitigation is a hard challenge. However, if early warning signals were to show an upcoming tipping point with catastrophic damage[1-7], theory and experiments suggest this could simplify collective action to reduce greenhouse gas emissions[8-11]. At the actual threshold, no country would have a free-ride incentive to increase emissions over the tipping point, but it remains for countries to negotiate their emission levels to reach these agreements. We model agents bargaining for emission levels using strategic reasoning[12,13] to predict emission bids by others and ask how this affects the possibility of reaching agreements that avoid catastrophic damage. It is known that policy elites often use a higher degree of strategic reasoning[13,14], and in our model this increases the risk for climate catastrophe. Moreover, some forms of higher strategic reasoning make agreements to reduce greenhouse gases unstable. We use empirically informed levels of strategic reasoning when simulating the model.

Game theoretic equilibrium analysis[9,10] shows that agents have incentives to coordinate[15] on agreements to avoid sufficiently severe catastrophic damage, but it is of importance to understand how negotiating agents starting from heterogeneous[16,17] bargaining positions can reach and form a basis for agreements in the first place[18]. Simulations of bargaining over emission levels suggest that agreement equilibria can be reached by agents using simple adaptive learning rules[9] comparable to fictitious play[19,20]. Experiments demonstrate that communicating groups of students tend to reach agreements with greater frequency when the threshold level is well-known compared with when there is uncertainty about the threshold level[8,11,21].

However, student experiments and simulations of adaptive agents may miss a critical aspect of the coordination problem. If policy elites are more skilled at strategic bargaining[13,14] we may expect a higher degree of strategic reasoning in international climate negotiations, which depend on experienced individual negotiators[22]. New survey experiments show that elite policymakers often use a higher level of strategic reasoning than students[13,14]. This could be relevant for our theoretical understanding about the coordination problem and can suggest a mechanism at play in climate negotiations[23].

Strategic reasoning which predicts others' strategies can have advantages from the perspective of an individual, but the overall effect on equilibrium selection[24] in bargaining is less clear. The purpose of our work is to model higher-order strategic reasoning to study if and how the level of strategic reasoning affects the possibility for bargaining players to reach agreement and avoid catastrophic damage.

We extend an n-player game theoretic model previously used to study emissions bargaining[9] and assume players represent identically sized countries in line with previous work on climate negotiations[8,11,21]. The model of the bargaining process has both cooperative and business-as-usual Nash equilibria and includes a finite population of n players indexed by the integers $1 \dots n$. Time is discrete in steps of 1 and the strategy of player i in round $r > 0$ is represented by a real number emissions bid $d_i(r) \in [0, 1]$. The players bargain over emission levels with simultaneous offers in a finite sequence of rounds. In the first round, $r = 1$, each player i chooses an initial bid $d_i(1)$ by a mechanism to be specified (see simulations below and Methods), and in the second round and onwards, $r \geq 2$, players simultaneously place bids to maximize emission payoff depending on the beliefs about other players' bids. A player's payoff is given by their emission level if there is no catastrophic damage, but if average emissions in the population exceed the threshold $T \in [0, 1]$ and total emissions exceed nT, players suffer from catastrophic damage that sharply reduces payoffs to fraction δ of emissions as in equation (1). The damage parameter δ thus also specifies the minimum guaranteed payoff when a player chooses a business-as-usual strategy with maximum emissions 1. Thus, total global emissions must be reduced to nT to avoid catastrophic damage, but a rational player will never choose an emission level below δ, because if there is less room than δ before passing the threshold, business-as-usual would be rational. For typical δ and T, when emissions total nT, every such agreement is a Pareto-efficient Nash equilibrium avoiding catastrophic damage (see Methods).

To model strategic reasoning in players, we use a level-k (Lk) model of players related to the behavioural economics and game theory literature[12,25–28]. Level-k models specify players on different levels of strategic reasoning. The L0 specification represents an L1 player's belief of other players, for example, that others play simply as they have in previous history, a belief to which L1 players best respond. Knowing this, L2 players predict others as L1 and make a rational response to this belief in maximizing payoff. Knowing this, L3 players can also correctly predict L2 players' bids and best respond to these beliefs (see Methods). Survey experiments with other games and an Lk model estimate elite policymakers to be predominantly L1 or L2 and higher, typically reasoning one step further than students estimated to predominantly match L0 (acting non-strategic) and L1 (ref. 13). These findings suggest that it is more likely to find higher levels of strategic reasoning in policy elites. Adaptive players in previous work[9] on climate bargaining correspond to a population of L1 players best responding to a L0 specification where others choose their average historical bids (see Methods and equation (6)).

Division of Physical Resource Theory, Department of Energy and Environment, Chalmers University of Technology, 412 96 Gothenburg, Sweden.
*e-mail: vive@chalmers.se

© 2016 Macmillan Publishers Limited. All rights reserved

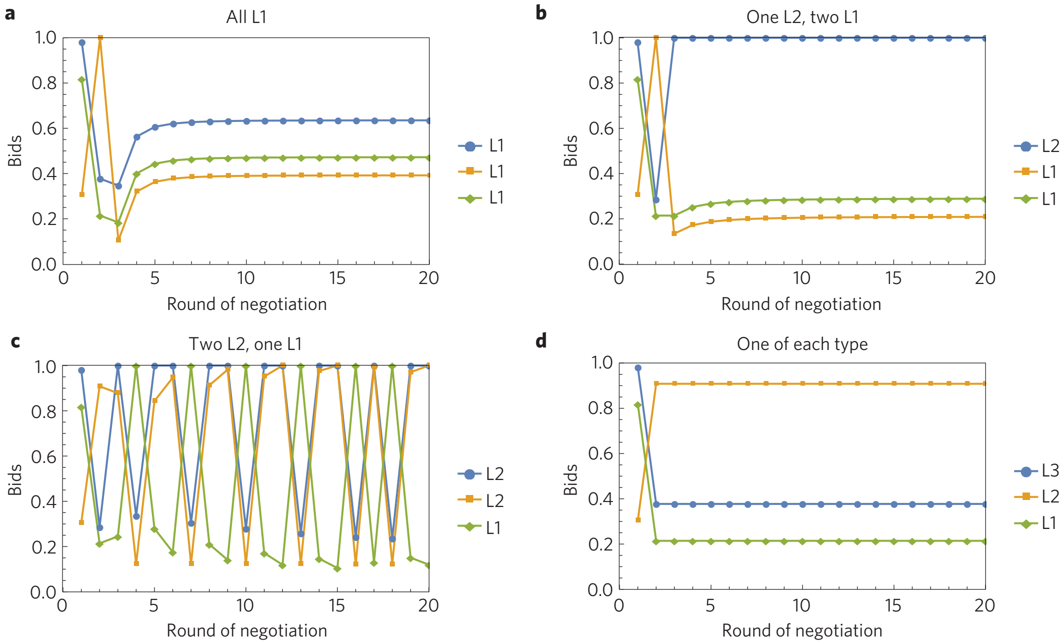

Figure 1 | Different levels of strategic reasoning lead to different bids for emission levels. Four bargaining scenarios, in each case the three players start with the same set of initial bids (0.980978, 0.30474, 0.817941) but with different levels of strategic reasoning. **a**, Represents the benchmark scenario of L1 players reflecting previous work[9]. **b–d**, One or more players with higher orders of strategic reasoning: one L2 and two L1 (**b**), two L2 and one L1 (**c**) and one each of L1, L2 and L3 (**d**). In the agreements reached in **b** and **d**, players with higher-order reasoning are better off than L1.

We study the effect of higher and heterogeneous levels of strategic reasoning on the outcomes of the climate bargaining model. We simulate climate bargaining by equations (2)–(7) and vary k of the Lk players in the range between 1 and 3 in line with previous findings for policy elites and where the levels typically are estimated. To illustrate the effect of higher-order reasoning, we study a baseline scenario with three players, a 50% reduction target and 90% damage if the threshold is passed ($n=3, T=0.5, \delta=0.1$), and compare with a population of L1 players (L0 given by equation (6)) without higher levels of reasoning.

In Fig. 1a, the baseline scenario shows an example in line with previous results[9] where L1 players gradually reach an agreement with aggregate emissions nT. Suppose that one of these players instead is L2 when starting with the same initial bids as in the previous case. Figure 1b shows how the L2 player with strategic reasoning has a clear advantage over the two L1 players. In round 2, the L2 player estimates the two L1 players' actions and predicts the maximum emission level that will not overshoot the target nT. This can be compared with the overshoot exceeding nT seen in round 2 of the baseline scenario in Fig. 1a. Moreover, in round 3, we see that the L2 player correctly predicts the possibility of choosing full emissions without overshooting the total emissions target, and the L1 players gradually reach an agreement in the following rounds because observing others makes them adapt their bids. The L2 player ends up better off in terms of payoff in the equilibrium agreement, because it quickly estimated that it could increase its emission level.

Because of bounded capability for strategic reasoning, players will predict other players incorrectly in some scenarios. Figure 1c shows two players as L2—that is, the L2 players best respond under the assumption that the others are both L1. An oscillating pattern of bids develops, and no agreement is reached. Higher levels of strategic reasoning can also allow a player to take into account that estimates by lower-level players can be wrong—for example, an L2 player incorrectly estimates an L3 player as L1. Figure 1d shows an example where an L3 player bargains with an L2 and an L1 player, respectively. The L3 player correctly estimates the two other bids by distinguishing between their levels. Here, bargaining will reach an equilibrium very quickly (explanation in Supplementary Information), but we see that the L2 player ends up better off as the L3 player rationally chooses to fill up the remaining emissions space (this advantage is typical, see Supplementary Information). Taken together, these examples suggest the possibility of a range of different outcomes depending on the levels of strategic reasoning among players in the population, with possible additional variability coming from parameters such as population size and initial bargaining positions.

To describe the overall effect of strategic reasoning, we vary levels of reasoning within a population of n players and simulate a large number of bargaining scenarios where initial bids are chosen uniformly at random in $[\delta, 1]$. Each simulated bargaining thus tests if any of the infinite number of possible agreements that avoid catastrophe are reached from a particular set of initial bids and expectations. The reduction target is set to 50% of business-as-usual ($T = 0.5$) and players' initial bids lie in the range of reducing emissions between 0 and 90% of business-as-usual ($\delta = 0.1$) and n is varied. Each outcome from the bargaining simulations is classified as either success (successfully reaching a cooperative equilibrium), failure (the business-as-usual equilibrium), or oscillating (no equilibrium is reached) after 100 rounds of negotiations (see Methods). Both reaching business-as-usual and continued oscillations are failures to reach agreement and no important effect was observed when increasing the number of rounds. Assessing the effect of strategic reasoning requires a benchmark for bargaining without higher-level reasoning; again, a natural reference is all players being L1.

Figure 2a,b shows how frequently simulations from 100,000 random starting positions reach the different types of bargaining outcomes. In Fig. 2a, we see that the baseline with only L1 players always has most success avoiding damage, except for with $n=3$ players. The success rates varies with reasoning levels and number of players, but we see that higher levels of strategic reasoning make fewer starting conditions lead to agreement. A striking difference is between a population of only L2 players compared to only L1 and L3,

© 2016 Macmillan Publishers Limited. All rights reserved

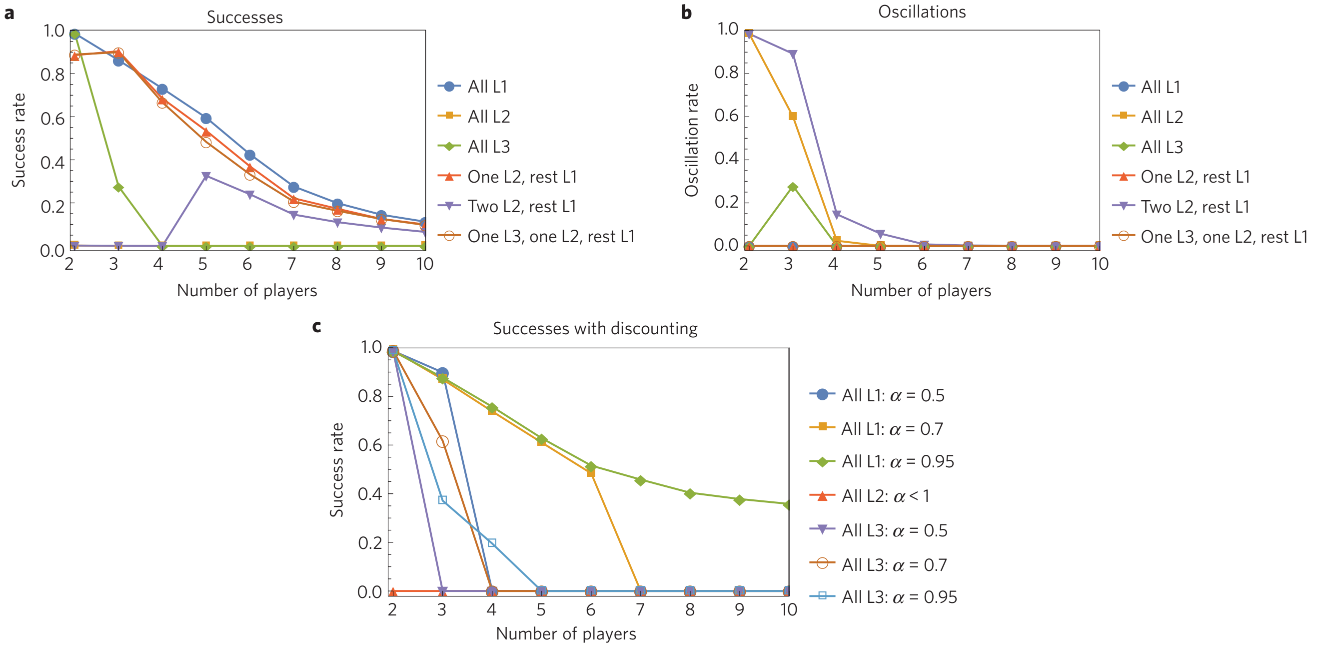

Figure 2 | The overall effect of strategic reasoning on bargaining outcomes when varying the number of players *n* and levels of strategic reasoning L1, L2 and L3. Each data point is based on negotiations starting from 100,000 random initial bargaining positions with maximum 100 rounds. **a**, Success rates in reaching agreements that avoid catastrophe. The L2 population always has zero success rate. **b**, Oscillation rates describing negotiations that do not settle to an equilibrium. **c**, Success rates for players with L0 specification given by exponential discounting. In all cases, $\alpha = 0$ leads to no success.

respectively, as the success rates for the L2 players are zero. Figure 2b shows that for smaller n the L2 players develop an oscillating pattern, but for larger n the L2 populations always reach business-as-usual. Reaching agreements for these populations seems difficult for two reasons. First, simulations show that L2 populations often quickly reach business-as-usual (see Supplementary Information). Second, for L2 populations agreements are also unstable (Proposition 1, Supplementary Information) and small perturbations lead away from the equilibrium. This model is obviously a number of steps away from real climate negotiations, but careful interpretation is that highly simplified negotiations exist where higher-order reasoning among strategically sophisticated negotiators make agreements to avoid catastrophe unstable. This suggests that strategic reasoning could be a roadblock to stay in agreements, unless small deviations are dampened.

A natural question is whether the overall effect of strategic reasoning carries over to different forms of L0 specification. For level-k models, the L0 specification can have important effects on how higher-order players form estimates[12]. Our level-k model is empirically informed by restricting players' reasoning levels in the range of previous estimates, but we examine the sensitivity of our results to the time-averaging L0 specification in equation (6) by considering an exponential smoothing[29] alternative defined by equation (7). This represents the idea that any degree of discounting the previous observations could be possible in L1 players' estimates of other players' emission bids. Consequently, more recent observations could be given more weight in the L0 specification (see Methods). We study the effect of discounting given by parameter $\alpha \in [0, 1]$ in equation (7) on bargaining outcomes.

Figure 2c shows outcomes from simulations where we varied discounting, strategic reasoning and number of negotiating players. Overall, higher-order reasoning makes fewer starting positions reach successful agreements. We again see the difficulty an L2 population has in reaching agreements, and for L1, L2 and L3 populations without sufficient weight on historical observations, agreement equilibria are unstable (Propositions 2–4, Supplementary Information). This sensitivity of agreement stability to what information is used to estimate others suggests once more it can be harder to stay in agreements unless small deviations are dampened.

One way to increase the success rate of finding agreement that was already found in previous work[9] is to restrict the range of initial bids in the first round so that negotiations start from less extreme bargaining positions closer to agreement. This measure has a similar effect also for players with higher levels of strategic reasoning (see Supplementary Information).

Taken together, strategic reasoning typically increases the risk of climate catastrophe in our model. This suggests that with potential catastrophic damage reaching emissions agreements is harder than what has been previously thought. When using simple game theoretic models for climate policy analysis we should strive to incorporate characteristics known to be likely for the negotiating agents—and strategic reasoning is one of them. The capability for strategic reasoning can thus seem useful for individuals but may also make it harder to cooperate to manage at least this global catastrophic risk[30]. Incorporating scientific findings about human behaviour as typically heterogeneous, out of equilibrium, and not always perfectly rational into our models can give new insights about the challenge to agree.

Methods

Methods and any associated references are available in the online version of the paper.

Received 14 August 2015; accepted 5 October 2015; published online 2 November 2015

References

1. Alley, R. B. *et al*. Abrupt climate change. *Science* **299**, 2005–2010 (2003).
2. Lenton, T. M. *et al*. Tipping elements in the Earth's climate system. *Proc. Natl Acad. Sci. USA* **105**, 1786–1793 (2008).
3. Scheffer, M. *et al*. Early-warning signals for critical transitions. *Nature* **461**, 53–59 (2009).
4. Lenton, T. M. Early warning of climate tipping points. *Nature Clim. Change* **1**, 201–209 (2011).

© 2016 Macmillan Publishers Limited. All rights reserved

5. Barnosky, A. D. *et al*. Approaching a state shift in Earth's biosphere. *Nature* **486**, 52–58 (2012).
6. Scheffer, M. *et al*. Anticipating critical transitions. *Science* **338**, 344–348 (2012).
7. Lenton, T., Livina, V., Dakos, V., Van Nes, E. & Scheffer, M. Early warning of climate tipping points from critical slowing down: Comparing methods to improve robustness. *Phil. Trans. R. Soc. A* **370**, 1185–1204 (2012).
8. Barrett, S. & Dannenberg, A. Sensitivity of collective action to uncertainty about climate tipping points. *Nature Clim. Change* **4**, 36–39 (2014).
9. Smead, R., Sandler, R. L., Forber, P. & Basl, J. A bargaining game analysis of international climate negotiations. *Nature Clim. Change* **4**, 442–445 (2014).
10. Barrett, S. Climate treaties and approaching catastrophes. *J. Environ. Econ. Manage.* **66**, 235–250 (2013).
11. Barrett, S. & Dannenberg, A. Climate negotiations under scientific uncertainty. *Proc. Natl Acad. Sci. USA* **109**, 17372–17376 (2012).
12. Crawford, V. P., Costa-Gomes, M. A. & Iriberri, N. Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *J. Econ. Lit.* **51**, 5–62 (2013).
13. Hafner-Burton, E. M., LeVeck, B. L., Victor, D. G. & Fowler, J. H. Decision maker preferences for international legal cooperation. *Int. Org.* **68**, 845–876 (2014).
14. Hafner-Burton, E. M., Hughes, D. A. & Victor, D. G. The cognitive revolution and the political psychology of elite decision making. *Perspect. Polit.* **11**, 368–386 (2013).
15. Stein, A. A. Coordination and collaboration: Regimes in an anarchic world. *Int. Org.* **36**, 299–324 (1982).
16. Vasconcelos, V. V., Santos, F. C., Pacheco, J. M. & Levin, S. A. Climate policies under wealth inequality. *Proc. Natl Acad. Sci. USA* **111**, 2212–2216 (2014).
17. Weikard, H.-P. & Dellink, R. Sticks and carrots for the design of international climate agreements with renegotiations. *Ann. Oper. Res.* **220**, 49–68 (2014).
18. Barrett, S. The theory of international environmental agreements. *Handbook Environ. Econ.* **3**, 1457–1516 (2005).
19. Fudenberg, D. & Levine, D. Learning in games. *Eur. Econ. Rev.* **42**, 631–639 (1998).
20. Young, H. P. *Strategic Learning and its Limits* (Oxford Univ. Press, 2004).
21. Milinski, M., Sommerfeld, R. D., Krambeck, H.-J., Reed, F. A. & Marotzke, J. The collective-risk social dilemma and the prevention of simulated dangerous climate change. *Proc. Natl Acad. Sci. USA* **105**, 2291–2294 (2008).
22. Kjellén, B. *A New Diplomacy for Sustainable Development: The Challenge of Global Change* (Routledge, 2014).
23. Victor, D. G. Copenhagen II or something new. *Nature Clim. Change* **4**, 853–855 (2014).
24. Samuelson, L. *Evolutionary Games and Equilibrium Selection* Vol. 1 (MIT Press, 1998).
25. Nagel, R. Unraveling in guessing games: An experimental study. *Am. Econ. Rev.* **85**, 1313–1326 (1995).
26. Stahl, D. O. Boundedly rational rule learning in a guessing game. *Games Econ. Behav.* **16**, 303–330 (1996).
27. Duffy, J. & Nagel, R. On the robustness of behaviour in experimental 'beauty contest' games. *Econ. J.* **107**, 1684–1700 (1997).
28. Shapiro, D., Shi, X. & Zillante, A. Level-k reasoning in a generalized beauty contest. *Games Econ. Behav.* **86**, 308–329 (2014).
29. Gardner, E. S. Exponential smoothing: The state of the art–part II. *Int. J. Forecast.* **22**, 637–666 (2006).
30. Bostrom, N. & Cirkovic, M. M. *Global Catastrophic Risks* (Oxford Univ. Press, 2008).

Acknowledgements

V.V., D.J.A.J. and K.L. want to thank the Swedish Energy Agency and the EU-FP7 project Mathemacs for financial support. We thank R. Smead, R.L. Sandler, P. Forber and J. Basl for providing detailed information about their Figure 1 starting values to facilitate comparison.

Author contributions

V.V. designed, performed the research and wrote the paper with contributions from D.J.A.J. and K.L.

Additional information

Supplementary information is available in the online version of the paper. Reprints and permissions information is available online at www.nature.com/reprints. Correspondence and requests for materials should be addressed to V.V.

Competing financial interests

The authors declare no competing financial interests.

© 2016 Macmillan Publishers Limited. All rights reserved

Methods

We extend a previous climate bargaining model with strategic reasoning. Let d_{-i} be the set of other players' emissions from the perspective of player i. Player i's payoff π_i depends on the player's emissions $d_i \in [\delta, 1]$ and others' emissions

$$\pi_i(d_i, d_{-i}) = \begin{cases} d_i & \text{if } \sum_{j=1}^{n} d_j \leq nT \\ \delta d_i & \text{otherwise} \end{cases} \tag{1}$$

where keeping emissions at or below nT avoids catastrophe. With $T > \delta$, this model[9] has a business-as-usual Nash equilibrium when $n > (1-T)/(1-\delta)$, and an infinite number of Pareto-efficient Nash equilibria determined by $\sum_i d_i = nT$ and $d_i \geq \delta$. The game is information imperfect, so in round r each player i knows the full history but lacks information about the current bids $d_{-i}(r)$ from other players and estimates these using strategic reasoning. Write $\tau_{-i}(r) = \sum_{j \neq i} d_j(r)$ for the total emissions by others in round r and $\hat{\tau}_{-i}(r)$ for player i's estimate of the relevant aggregate $\tau_{-i}(r)$.

Then, $D(i,r) = nT - \hat{\tau}_{-i}(r)$ is the estimated emissions room for player i before exceeding the threshold at nT. Best response to this belief and equation (1) is

$$d_i(r, \hat{\tau}_{-i}(r)) = \begin{cases} \min(1, D(i,r)) & \text{if } D(i,r) \geq \delta \\ 1 & \text{otherwise} \end{cases} \tag{2}$$

reflecting that a player switches to business-as-usual if expecting emissions room before catastrophic damage to be less than δ. We introduce strategic reasoning of level-k (Lk) players by different estimates $\hat{\tau}^k_{-i}(r)$ of others' bids $\tau_{-i}(r)$ in the population. (For interpretation and comparison of our model to level-k models, see Supplementary Information.) L1 players are defined by their best response to the estimate given by the L0 specification $\hat{\tau}^0_{-i}(r)$

$$d^1_i(r) = d_i(r, \hat{\tau}^0_{-i}(r)) \tag{3}$$

using the L0 estimate to predict the relevant aggregate from $n-1$ other players. Similarly, L2 players are defined by

$$d^2_i(r) = d_i(r, \hat{\tau}^1_{-i}(r)) \tag{4}$$

where $\hat{\tau}^1_{-i}(r) = \sum_{j \neq i} d^1_j(r)$ estimates others as L1 decision makers.

L3 players are defined by best response to the higher-order estimate $\hat{\tau}^2_{-i}(r)$

$$d^3_i(r) = d_i(r, \hat{\tau}^2_{-i}(r)) \tag{5}$$

where $\hat{\tau}^2_{-i}(r) = \sum_{j \in L1} d^1_j(r) + \sum_{j \in L2 \cup L3 \setminus \{i\}} d^2_j(r)$. Here L1, L2 and L3 denote the sets on $\{1 \ldots n\}$ for the corresponding level-k players. L3 players estimate other L3 players as L2, where we have assumed they distinguish between L2 and L1 players, and in general an Lk player estimates remaining Li players with $i \geq k$ as L$(k-1)$ in line with level-k models[12]. We consider the L0 specification to be on the form

$$\hat{\tau}^0_{-i}(r) = \frac{1}{r} \sum_{k=0}^{r-1} \tau_{-i}(k) \tag{6}$$

including a pre-negotiation state[9] $d_i(0) = d_i(1)$ to describe initial expectations when bargaining begins. Equation (6) represents taking the historical average over $n-1$ other players to estimate the relevant aggregate $\tau_{-i}(r)$. If we take all players to be L1, equations (3) and (6) directly describe previous work[9], whereas equations (4) and (5) describe higher levels of strategic reasoning. For a different way of including previous bids in the L0 specification and for sensitivity analysis, we consider exponential discounting which over time gives more recent observations more weight

$$\hat{\tau}^0_{-i}(r) = \alpha \hat{\tau}^0_{-i}(r-1) + (1-\alpha)\tau_{-i}(r-1) \tag{7}$$

with parameter $\alpha \in [0,1]$ specifying how much the L0 specifications $\hat{\tau}^0_{-i}(r-1)$ in the previous round are weighted against the last observations. This represents players discounting previous observations at rate α: with $\alpha = 0$, a player's estimate always corresponds to the last round, and with $\alpha = 1$, observations do not change the estimate at all. This can be viewed as a form of learning with exponential smoothing[29] (see Supplementary Information for more details).

Simulation details. We structure the following to facilitate comparison with Smead and colleagues[9].

Initialization. For each simulation, we initialize n players and, except for where noted, their initial bids are taken uniform at random in their strategy set $[d_{\min}, d_{\max}]$. When there are no restrictions on their initial demands, $[d_{\min}, d_{\max}] = [\delta, 1]$, and when there are restrictions, we have that $[d_{\min}, d_{\max}] = [\delta + R, 1 - R]$, where R has been taken to satisfy $\delta + R < 1 - R$. Each player is given a level of reasoning and is classified either as L1, L2 or L3 at random, unless where noted. Let $r = 2$ to denote that the bids in the first round and that the pre-negotiation state $d^0_i(0)$ are fixed comparable to previous work[9].

Each round of the simulation. In each round $r \geq 2$ of the simulation, for each player i, we compute either equation (6) or equation (7), depending on which of the two L0 specifications is considered. Write $\hat{\tau}^0_{-i}(r)$ for the result below.

Then, using $\hat{\tau}^0_{-i}(r)$ for each i, for L1 players we compute equation (3), for L2 players we compute equation (4), and for L3 players we compute equation (5). Let the result for the respective players be the bids $d_1(r), d_2(r), \ldots, d_n(r)$ in round r. Then, we increase r by 1, advancing the simulation to next round.

Stopping the simulation. If $r = 100$, the simulation is stopped and the outcome is classified.

Classifying the results. From $r = 100$ we perform an additional simulation of 50 rounds, where r takes steps from 101 to 150 to classify the result. The following three criteria were used: *Success* (Reaching agreement): if for each $r' > r$ and $r' \in [101, 150]$ the population total and each player is within 0.1% of nT and the same holds for each player's bid. This holds if $|nT - \sum_i d_i(r')| < 0.001$ and $|d_i(r') - d_i(100)| < 0.001$ for each player i. *Failure* (business-as-usual): if, as for success, the population is within 0.1% of business-as-usual play emissions (all 1, $\sum_i d_i(r') = n$). *Oscillating* (no signs of settling to an equilibrium): if an outcome did not fit the criteria of success or failure above, we call it oscillating.

© 2016 Macmillan Publishers Limited. All rights reserved