Extended abstract:

Game theory has had great difficulty dealing convincingly with an important family of games whose equilibria may be found by a long backward induction. Exemplars of this problematic but central family are the finitely repeated prisoner's dilemma [Luce Raiffa 57] and Rosenthal's centipede [Rosenthal 81], [Kreps 90]; Basu's traveler's dilemma [Basu 94] and imperfect price competition games [Capra et al 02] are related examples. On the one hand, the theory of Nash equilibrium clearly predicts one sort of behavior in these games, which we could call a race to the bottom; on the other hand, this behavior seems unreasonable, even for rational, self-interested agents, an intuition reflected in a large number of experiments, including those with subjects well aware of the Nash equilibrium [Becker et al 05]. Many (overlapping) attempts have been made to address this seeming paradox: for example, the introduction of a small chance of irrationality of a carefully chosen type [Kreps et al 82], the introduction of limits on the reasoning or computational power [Neyman 97] of agents, or a cost for the use of complex strategies, the weakening of maximizing behavior to near-maximizing behavior [Radner 80], [Simon 55], and the replacement of maximizing behavior with a sort of stochastic or smoothed maximization [McKelvey Palfrey 95]. There are difficulties with each of these approaches. For example, the first is open to the objection that alternative choices of the seeding irrationality lead to a wide variety of outcomes [Fudenberg Maskin 86], the second seems not to apply well to games with simple structures like centipede or traveler's dilemma, especially with sophisticated agents, the third typically lacks specificity in its predictions, and the fourth depends on the choice both of a type of smoothing function and a diffusion parameter.

Our approach is closest in spirit to the third and fourth lines of attack, though we believe it possesses significant advantages. While it is much simpler than the stochastic framework of quantal response equilibria, it involves an often radical refinement of the set of ε-equilibria despite the introduction of no further parameters. This allows for quite sharp predictions in the case of centipede and traveler's dilemma which are in striking agreement with experiments and intuition.

We were motivated originally by the much studied paradox of the finitely repeated prisoner's dilemma, which is more than half a century old. Considering a prisoner's dilemma repeated 100 times, where every Nash equilibrium leads both players to play tough on every round, Luce and Raiffa [Luce Raiffa 57] state that they would not play to a Nash equilibrium. In fact, if strategies were restricted to those which play nice before some round k, from 1 to 100, as long as the opponent also plays nice, and after a tough play by the opponent or the arrival of round k play tough until the end, they write that they would probably play a strategy k, where k "is some number in the nineties." We are able to vindicate their intuition for the restricted strategy game; the unrestricted game is still beyond our grasp. In the same way, we resolve the paradox of the centipede game and the traveler's dilemma. We also address some limiting instances of stag hunt, the prototypical assurance game whose history dates to Rousseau [Skyrms 04], [Battalio et al 01], as well as some other relevant

examples from the literature.

We introduce here a small circle of closely related solution concepts for games in strategic form centered on the notions of ε-dominance and ε-robustness. Our aim is to expand the normative and positive scope of noncooperative game theory with the simplest possible tools.

A Nash equilibrium requires zero regret from each agent if he has correctly anticipated others' strategies, but allows massive regret if another's strategy is unforeseen. This makes Nash equilibria precariously dependent on very strong assumptions. We instead require small regret from each agent if he predicts correctly, but also impose some robustness on his strategy, i.e. seek to limit his regret if his prediction is incorrect. This turns out to be surprisingly fruitful.