



Coalitional stochastic stability

Jonathan Newton

School of Economics, University of Sydney, Australia

ARTICLE INFO

Article history:

Received 13 April 2010

Available online 3 March 2012

JEL classification:

C71

C72

C73

Keywords:

Stochastic stability

Learning

Coalition

Lexicographic

Contract

ABSTRACT

This paper takes the idea of coalitional behavior – groups of people occasionally acting together to their mutual benefit – and incorporates it into the framework of evolutionary game theory that underpins the social learning literature. An equilibrium selection criterion is defined which we call *coalitional stochastic stability* (CSS). This differs from existing work on stochastic stability in that profitable coalitional deviations are given greater importance than unprofitable single player deviations. A general characterization of CSS is given together with more detailed characterizations for specific classes of games. Applications include contracting, asymmetric social norms and collusive price setting, the latter of which is shown in some circumstances to facilitate competitive outcomes.

© 2012 Elsevier Inc. All rights reserved.

1. Introduction

In their seminal papers in evolutionary game theory Foster and Young (1990) and Kandori et al. (1993) introduce the idea of *stochastic stability*: a method of equilibrium selection which assesses the robustness of equilibria by measuring their resilience to random errors in players' actions. The model of Foster and Young features a continuous dynamic whereas Kandori et al. look at discrete dynamics and give a graphical approach to finding the stochastically stable states. Young (1993) applies these ideas to the evolution of social norms in a discrete time dynamic process which he calls adaptive play. Under adaptive play, players repeatedly play an n -player game Γ . We note that Young interprets adaptive play as modeling situations where each player is actually a representative agent picked at random from some population of similar agents. This interpretation is also a valid possibility for this paper, but for the sake of clarity we stick to using the term 'player' to describe a position in a game, whether it is the same agent repeatedly playing or various agents plucked randomly with replacement from some underlying population. Players follow a process whereby they play best responses to a distribution over the actions played by the other players, where the distribution is determined by sampling s actions from the previous m actions played by the other players. This defines a Markov process where the states of the process are defined by the actions taken by each of the players in the previous m periods. If s is small enough relative to m and the game is *weakly acyclic* the game converges almost surely to a *convention* where each player has played a strategy from a pure strategy Nash equilibrium of Γ for as long as anybody can remember. Young introduces random shocks to the system which can be interpreted as random mistakes made by players in implementing their strategies. As long as there is a positive probability of every combination of strategies being played by mistake the perturbed system then defines an aperiodic and irreducible Markov process. Young shows that as the probability of a random shock ϵ approaches zero, the system spends almost all of its time in a subset of the conventions of the games. He calls such conventions *stochastically stable*.

E-mail address: jonathan.newton@sydney.edu.au.

This paper proposes *coalitional stochastic stability*: an alternative to random error driven stochastic stability for situations where coalitional behavior is more likely than random errors. Coalitional stochastic stability uses the possibility of deviations by groups of players as an equilibrium selection device in evolutionary models of social learning. This paper describes, justifies and illustrates the uses of this innovation, giving a general characterization of CSS states as well as more detailed characterizations for two-player games, coordination games and supermodular games. It is shown that in games of contracting the Nash bargaining solution is selected by CSS, in contrast to a previous result of Young (1998a) and in accordance with the predictions of Naidu et al. (2010), the effects of whose paper arise endogenously in the current paper. Applications are given throughout the paper, illustrating how CSS can be used to give theories of the persistence and direction of inequalitarian social norms (Section 4.3) and how collusive behavior between firms can promote (as well as inhibit) price competition in Bertrand models (Section 4.8).

Section 2 of this paper summarizes the two areas of related literature and gives the motivation and contribution of this paper in more detail. Section 3 describes the model. Section 4 includes the fundamental characterization results of the paper and several examples and applications. Section 5 concludes. Formal proofs are given in Appendix A.

2. Related literature

There are two strands of literature which this paper bridges. Below I give a brief summary of both of them, followed by a description of the motivation and contribution of this paper.

2.1. Coalitional behavior

There exists a large literature in cooperative game theory on the behavior of coalitions. For a survey the reader is referred to Peleg and Sudholter (2003). Aumann (1959) gives the concept of a ‘strong equilibrium’ – an equilibrium where no subset of players would want to agree to change their profile of strategies to another profile, holding the strategies of all players not in that subset fixed. This equilibrium concept can be argued to correspond to situations where coordination between any subset of players is possible without players outside the subset being aware of it. As the name suggests this is a very strong equilibrium notion and often will not exist. The concept can be weakened to that of k -strong equilibrium where only coalitions of size k or lower have to have their incentive constraints satisfied but still existence is not guaranteed.¹ Bernheim et al. (1987) attempt to address the issue of robustness to coalitional deviations through their concept of coalition proof equilibrium, the idea of which is that equilibria need to be robust to a set of players deviating only if that set of players is itself robust to any further deviations by subsets of its constituent players. Bernheim et al. argue that this equilibrium concept can be understood intuitively to lead to outcomes which could be reached if all players seated in a room reached an agreement, following which the players leave the room one by one, and no matter in what order they leave the room there will never be a subset of players remaining in the room who would agree to play differently to what was agreed with all players present. Konishi and Ray (2003) look at the issue of coalition formation in a dynamic setting with farsighted agents, showing that if characteristic functions are incorporated into the rules governing the dynamic process then they can choose such a process which always selects payoffs in the core of the underlying game if the core is a singleton. Ambrus (2009) defines and analyzes a concept of coalitional rationalizability: the idea that subsets of players will refrain from playing certain strategies if it is in their interests to do so. Luo and Yang (2009) extend this concept to situations where players use Bayesian updating to calculate expected payoffs.

2.2. Stochastic stability

Whereas previous concepts of equilibrium stability such as asymptotic stability or evolutionary stable strategies (Smith and Price, 1973) focus on robustness to single errors (mutations) in strategies, Young (1993) and Kandori et al. (1993) introduce the possibility of multiple ‘random errors’ in strategies chosen and show that although there may be several stationary states in a dynamic process, some of them may be more robust to such errors than others, and that if the probability of a random error in the very long run becomes very small, then the state which is most robust to such errors will be observed almost all of the time. Kandori et al. (1993) and Young (1993) predict that in 2×2 games when there are two strict Nash equilibria then the risk dominant equilibrium will be selected. Ellison (2000) provides some further useful characterization results. Bergin and Lipman (1996) prove a kind of folk theorem for stochastic stability, that is they show that any stable state of the unperturbed dynamic process can be selected with appropriately chosen state-dependent mutation rates. van Damme and Weibull (2002) recover some of the predictive power of the theory by assuming that avoiding mistakes is costly to players. They endogenize the random error probabilities so that players will pay more to avoid making mistakes which are more costly to them and give conditions under which the probabilities of random errors by any player at any state are of the same order of magnitude and remain so as limits are taken. They show that under these conditions the results of Young are recovered. Young (1998a) shows that under his uniform error stochastic stability process there is a preference for fairness in contracts agreed between two players and that the contract selected corresponds

¹ Nash equilibrium is a special case of k -strong equilibrium where $k = 1$.

to the Kalai–Smorodinsky bargaining solution. Naidu et al. (2010) analyze a model of contracting where movement from one Pareto efficient contract to another is only caused by errors on the part of the player who stands to gain from the move. They justify this by invoking a level of foresight on the part of the agents, who know that there is a better contract available and thus ‘intentionally’ make mistakes so as to lead to a better conventional contract for themselves.

2.3. Motivation of this paper

Instead of perturbing rationality to obtain predictions in games, this paper focuses on perturbing individuality. We look at environments where from time to time players may meet and agree to jointly coordinate their actions. We develop a stochastic stability notion which we call *coalitional stochastic stability* (CSS) which is based on resilience to coalitional behavior, unlike traditional stochastic stability notions which are based on resilience to random errors.

It is one thing to analyze the behavior of a dynamic where coalitional behavior is possible. It is another thing to analyze the predictions of a model where coalitional behavior becomes infinitesimally likely in the limit. Aside from pure academic interest there are several reasons why we should be interested. Firstly, people deal with a lot of games in their everyday lives and the amount of time they devote to any given one is by necessity limited. It is not unreasonable to think that it is quite a rare event that two or more people get together and discuss any particular aspect of their lives and credibly agree to make the necessary changes in strategy to better their outcomes. It is also not unreasonable to think that the more people required to agree to changes in behavior, the harder these changes are to accomplish. Secondly, for the results of this paper it is not necessary that individual strategic switching occur with positive probability in the limit, merely that it is much more likely than coalitional strategic switching. Thus, any type of strategic change can be viewed as a rare event, a fair assumption given the stasis observed in people’s behavior in most aspects of their lives. Thirdly, limiting results approximate what would be observed in the presence of strictly positive amounts of coalitional activity and provide insight into the strategic dynamics that coalitional activity induces. We note that the behavioral implications of the most interesting propositions and examples in the paper do not rely on limiting arguments.

3. Model

This paper shall follow closely the methods and notation of Young (1993).² Take a game Γ with set of players N , $|N| = n$, finite strategy sets X_1, \dots, X_n ; $X = \prod X_i$; and payoffs given by $\pi_i: X \rightarrow \mathbb{R}$. The action taken by player i at time t is denoted x_i^t . The action profile played at time t is denoted $x^t = (x_1^t, \dots, x_n^t)$. The *state* of the system is given by the actions played in the last m periods and is denoted $h^t = (x^{t-m+1}, \dots, x^t)$. The higher the value of m the longer the memory of the players. X^m denotes the set of all possible states. The system is taken to start at an arbitrary state: it is assumed at least t periods have already elapsed since the beginning of time.

We define a Markov process P_0 on X^m as follows:

- From state h^t player i draws a random sample of s actions out of the previous m actions taken by each of the other players. These $n - 1$ samples are drawn independently. The sample distribution of j ’s actions in i ’s sample is denoted \hat{p}_{ij}^t and we write $\prod_{j \neq i} \hat{p}_{ij}^t = \hat{p}_{-i}^t$
- Player i plays a best response to \hat{p}_{-i}^t . We denote the set of such best responses by $B_i(\hat{p}_{-i}^t)$. If there exist tied best responses they are played with equal probability. The actions all players take define x^{t+1} and thus the next state h^{t+1} .

A *convention* is defined as a state $h^t = (x^{t-m+1}, \dots, x^t)$ where $x^{t-m+1} = x^{t-m+2} = \dots = x^t = x^*$ and x^* is a strict Nash equilibrium of the underlying game Γ . It is clear that under the process P_0 once a convention is reached it will be sustained forever. However, Young defines a perturbed version of the process P_ϵ where with probability $1 - \epsilon$ a player plays a best response as per P_0 . With probability ϵ he instead makes an ‘error’ and plays a random action from a distribution with full support over his possible actions. In the limit as $\epsilon \rightarrow 0$ only some conventions are played with positive probability in the long run – they are *stochastically stable*.

3.1. Coalitional stochastic stability

In this section we introduce *coalitional stochastic stability*. Like standard stochastic stability it is a concept based on limits and as such shares the benefits of its sharp predictive precision. CSS primarily concerns coalitional deviations but also includes random errors as a technical tool to ensure irreducibility of the Markov process. There are many games where this will make no difference to predictions, and where it does, CSS without random errors will give sharp predictions for large classes of states – precisely those classes of states which are closed under intentional coalitional behavior.

First define the following notation. \mathcal{P} is the set of all subsets of N . $\mathcal{P}_k \subset \mathcal{P}$ is the set of all subsets P_k of N such that $|P_k| = k$. Let F_k be a probability distribution over \mathcal{P}_k with full support. Given a state h^t , samples \hat{p}_{-i}^t , $a_i \in B_i(\hat{p}_{-i}^t) \forall i$, and a set $Q \subseteq N$, let:

² See also Young (1998a, 1998b).

$$\begin{aligned}
 A_Q^{(1)}(h^t) &= \{x: x_i \in B_i(\hat{p}_{-i}^t) \forall i \notin Q\}, \\
 A_Q^{(2)}(h^t) &= \{x: \mathbf{E}_{\hat{p}_{-i}^t}[\pi_i|x_j, j \in Q] \geq \mathbf{E}_{\hat{p}_{-i}^t}[\pi_i|x_i = a_i] \forall i \in Q\}, \\
 A_Q^{(3)}(h^t) &= \{x: \nexists \tilde{x} \text{ s.t. } \mathbf{E}_{\hat{p}_{-i}^t}[\pi_i|\tilde{x}_j, j \in Q] \geq \mathbf{E}_{\hat{p}_{-i}^t}[\pi_i|x_j, j \in Q] \forall i \in Q, \\
 &\quad \mathbf{E}_{\hat{p}_{-i}^t}[\pi_i|\tilde{x}_j, j \in Q] > \mathbf{E}_{\hat{p}_{-i}^t}[\pi_i|x_j, j \in Q] \text{ for some } i \in Q\}, \\
 A_Q(h^t) &= A_Q^{(1)}(h^t) \cap A_Q^{(2)}(h^t) \cap A_Q^{(3)}(h^t)
 \end{aligned}$$

$A_Q(h^t)$ is thus the set of action profiles where (1) all players $i \notin Q$ play best responses to their sample distributions as in standard adaptive play, (2) the players in Q play actions such that the expected payoff of a player $j \in Q$ conditional on his knowing the actions of the other players in Q and given his sample distribution of the actions of players outside Q is higher than his expected payoff in standard adaptive play, and (3) the actions of players in Q are Pareto efficient for Q and can be regarded as a coalitional best response. Let $G(A_Q(h^t))$ define a probability distribution over $A_Q(h^t)$ with full support. Let \hat{H} be a probability distribution with full support over N and let H_i be distributions with full support over all possible actions of player i . Let b_2, \dots, b_{n+1} be positive constants.

Consider the following perturbed adaptive process P_ϵ with $\epsilon > 0$:

- With probability $1 - \sum_{i=2}^{n+1} \epsilon^{b_i}$ players follow the adaptive learning process as usual.
- With probability ϵ^{b_k} there is a Pareto superior deviation by $1 < k \leq n$ players. To be precise, a set of players P_k is selected according to F_k . All $i \in N \setminus P_k$ play best responses to their sample distributions as normal. The actions of $i \in P_k$ change so that payoffs for these players are weakly better under the new action profile than they would be under individual best responses: given that the state in the previous period was h^{t-1} , a new action profile x^t is selected from the distribution $G(A_{P_k}(h^{t-1}))$ and played in the current period. We call these deviations *coalitional deviations of order k*.
- With probability $\epsilon^{b_{n+1}}$ a random error occurs to the strategy of a randomly selected player. A player is selected according to \hat{H} and he plays an action determined by H_i . All other players play best responses to their sample distributions as normal.

As the process is irreducible it has a unique stationary distribution which I denote (à la Young, 1998b) as μ_ϵ . Let $\beta = \min_k \{\frac{b_{k+1}}{b_k}\}$. A state h is *coalitionally stochastically stable* if there exists $\hat{\beta}$ such that:

$$\beta > \hat{\beta} \implies \lim_{\epsilon \rightarrow 0} \mu_\epsilon(h) > 0.$$

3.2. What does this mean?

This definition ranks the different types of deviations in order of importance. Most important are profitable single player deviations, followed by profitable two player deviations and so on. Least important of all are random unprofitable deviations. In working out our CSS states we take in turn:

- Any type of intentional coalitional deviation to be infinitely more likely than unprofitable deviations.
- Intentional coalitional deviations involving fewer players to be infinitely more likely than intentional coalitional deviations involving more players.

So it is apparent that this concept puts a high value on rationality relative to the approach of Young (1993). Conversely, it puts lower emphasis on the independence of the players' actions. It can be understood to model a situation where the players in each position in the game are drawn each period from a population and, with some knowledge of how the game has been played in the past, best respond to how they expect their opponents to play, with the additional possibility that every so often groups of players will get together and adopt strategies to their mutual benefit. The ordering of the coalitional deviations can be understood to represent a situation where the cost of coalition formation increases in the size of coalitions or where small coalitional deviations are more easily replicable. This could be driven, for example, by smaller groups of players that have discovered a profitable deviation being more likely to be rematched in the same game positions in the future. Conversely, one can think of situations where news of large coalitional deviations is more likely to spread due to the larger number of players involved. Hwang (2011) even suggests that coalitional activity may be more easily enforced in larger groups. The theory in the paper can also be implemented in this reverse case. Occasionally throughout the remainder of the paper we will comment on how a different ordering of coalitional deviations would affect outcomes.

	L	M	N	O	R
a	6, 2	4, 1	30, 0	0, δ	0, δ
b	1, 8	3, 7	0, 0	0, δ	0, δ
c	0, δ	0, δ	0, 0	7, 3	1, 5
d	0, δ	0, δ	0, 0	8, 1	2, 6

Fig. 1. A two player strategic game. δ is assumed to be positive and close to zero.

4. Results and examples

4.1. Fundamental propositions

The following propositions establish the existence of CSS states, give a method for finding them, and link them to the Nash equilibrium concept.³ Define $\mathcal{R} := \mathbb{N}^n$, $r = (r_2, r_3, \dots, r_{n+1}) \in \mathcal{R}$, define addition componentwise for elements of \mathcal{R} , and define a reverse lexicographic ordering on \mathcal{R} . For $r^\dagger, r^\ddagger \in \mathcal{R}$: $r^\dagger \geq r^\ddagger \Leftrightarrow (r_i^\dagger < r_i^\ddagger \Rightarrow \exists j > i: r_j^\dagger > r_j^\ddagger)$. Define the resistance r of a path between two states by setting r_i equal to the number of coalitional deviations of order i on the path, r_{n+1} equal to the number of random errors on the path. Define the resistance $r(h_1, h_2) \in \mathcal{R}$ between states h_1, h_2 as the minimum resistance of any path from h_1 to h_2 . Denote the components of $r(h_1, h_2)$ by $r_i(h_1, h_2)$. Define a *spanning tree* rooted at h as a directed graph on the set of all states such that a unique path leads from every state to h . The cost of a tree is given by the sum of the resistances of its edges. The *stochastic potential* of a state h is defined as the cost of the lowest cost spanning tree rooted at h .

Proposition 1. *CSS states exist and are identical to the states with minimum stochastic potential.*

Proposition 2. *Under conditions which in Young (1993) guarantee the selection of a convention(s) by SS, CSS will also select a convention(s). These conditions are:*

- The underlying game Γ is weakly acyclic.
- The sample size s is sufficiently small relative to m .

Proposition 3. *If the set of conventions is partitioned into two sets C_1, C_2 , and there exists some k such that $\forall h_1 \in C_1, \forall h_2 \in C_2: r_k(h_1, h_2) > 0; \forall h_2 \in C_2: \exists h_1 \in C_1: \forall i \geq k: r_i(h_2, h_1) = 0$, then all CSS states are contained in C_1 .*

These propositions extend the results of Young (1993) and Ellison (2000) to the model of differentiated errors analyzed here. The lexicographic nature of $r(\dots)$ makes the application of Proposition 3 especially simple, as shall be seen in several examples. It is also useful to note that when $h_1 \in X^m$ is not in a recurrent class of P_0 , then there exists a member h_2 of a recurrent class such that $r(h_1, h_2) = 0$. This means that for the purpose of calculating the spanning trees of Proposition 1, transient states of P_0 can be ignored.

4.2. Example

Here we demonstrate how Proposition 1 can be used to find CSS states. δ is assumed to be positive and very close to zero and is included so that action N is strictly dominated for player 2. The 4×5 game in Fig. 1 has 2 strict Nash equilibria, aL and dR , each of which corresponds to a convention of our unperturbed dynamic. Under the standard random errors approach, convention aL will then be selected as it takes relatively few errors by player 2 where he chooses N (a strictly dominated action for him) before it becomes worthwhile for player 1 to play a in the hope of earning a very big payoff at aN . Once a is being played by player 1 it then becomes a best response for player 2 to play L and the convention aL is reached. Under CSS, from convention dR enough two player deviations to bM will allow aL to be reached, and from convention aL enough two player deviations to cO will allow dR to be reached. We now use deviations of order 2 to choose between conventions aL and dR . From dR the only possible two player coalitional deviation is to bM and $\frac{2}{3}$ of these are required before it is possible that player 1 sampling from player 2's actions sees player 2 playing M often enough for player 1 to judge it worth his while to play a . If player 1 then plays a for a while, player 2 can begin to play L and aL is reached. On the other hand, to move from aL to dR via plays of cO requires only $\frac{2\delta}{7}$ joint deviations before player 2 can switch to playing R followed by player 1 switching to d . So $r(aL, dR) = (\frac{2\delta}{7}, 0)$, $r(dR, aL) = (\frac{2}{3}, 0)$ and dR is selected by Proposition 1 as the unique coalitionally stochastically stable convention.

We have chosen between two conventions without either player engaging in behavior that is detrimental to his short term myopic best interest. Irrational, unintentional behavior is not always necessary in order for stochastic stability arguments to have bite.

³ See Appendix A for proof of these and all subsequent propositions.

	O	F
O	3, 3, 4	6, 0, 4
F	0, 6, 4	0, 0, 0
	O	

	O	F
O	0, 0, 0	0, 4, 6
F	4, 0, 6	2, 2, 6
	O	F

Fig. 2. Battle of the sexes with 2-1 sex ratio.

4.3. Example: Battle of the sex ratio

Fig. 2 illustrates an extended version of the battle of the sexes, the difference being that there are two players, the row and column players (RC players), who prefer opera (O) to football (F) and wish to coordinate with the matrix player (M). It is assumed that if they both coordinate with M they have an equal chance of being chosen. M prefers football to opera and wishes to coordinate with at least one of the other players. The example could represent any of a number of situations where one agent wishes to agree on a contract with either of two other agents.⁴ In a symmetric two player battle of the sexes both Nash equilibria are stochastically stable. The introduction of an additional opera-favoring player might be expected to give the football-favoring player some market power. In fact, under standard stochastic stability the opposite happens: for large enough sample sizes, the opera equilibrium becomes uniquely stochastically stable. There are two effects which cause this outcome. The first effect is that even when one RC player has mutated away from the current convention, there is still an RC player who remains playing that convention. This increases the expected payoffs to M from sticking with the existing convention and makes the resistances of transitions which are payoff improving for M higher than in the two-player game. The second effect is that lower payoffs for the RC players due to sharing lead to a greater willingness of individual RC players to experiment with different actions when they observe behavior by M that is not in keeping with the current convention. This effect is greater for transitions which decrease M's payoff than for transitions which improve his payoff.

This curious result is driven by the probabilities given to mutation patterns in each period. Within the framework of Young (1993) payoff improving deviations by a single player have a high probability of being realized. Payoff decreasing deviations by a single player have a probability of order ϵ of being realized. Deviations by two players such that either deviation on its own would be detrimental to the deviator's payoff have a probability of order ϵ^2 of being realized, *whether or not the deviations taken together offer a Pareto improvement to the deviators*. Thus as $\epsilon \rightarrow 0$, joint deviations by two players which offer payoff improvements become infinitely unlikely compared to unprofitable deviations by a single player. This is clearly not satisfactory for all situations, especially situations where coalitional behavior by groups of players is possible. In the battle of the sex ratio, if the current convention is for O to be played, then such a joint deviation could involve M and one of the RC players agreeing to play F. M would clearly benefit from this deviation and the RC player would also benefit as she would obtain the coordination payoff with probability 1 rather than with probability $\frac{1}{2}$.

In our model, profitable two player joint mutations occur with higher probability than unprofitable single player mutations. This selects the equilibrium where all players choose F.⁵ That is, when intentional coalitional behavior is given sufficient priority over unintentional single player deviations, the market power given to M by there being more than one RC player in the market is sufficient to take him to his favored equilibrium (nb. In the above game this happens by a kind of Bertrand argument whereby if M is not at his optimal contract M and an RC player deviate to a contract where they both do better in the short run and the other RC player follows so as to earn a non-zero payoff).

4.4. Efficiency and coordination

Two-player games are a special case when it comes to examining coalitional behavior because any coalitional move in a two-player game is by definition a move towards efficiency. In fact, if Γ is a two-player normal form game with multiple strict Nash equilibria then CSS will never select an equilibrium which is Pareto inferior to another equilibrium which is itself not weakly dominated by another cell of the payoff matrix.

Proposition 4. *When Γ is a weakly acyclic two-player game with a strict Nash equilibrium x^* which is both Pareto efficient and Pareto dominates all other strict Nash equilibria, then x^* is uniquely selected by CSS.*

What might be considered surprising is that Pareto dominance in the set of equilibria is not in itself sufficient to guarantee selection by CSS. It is possible that a non-equilibrium payoff vector could dominate the payoffs in a dominant equilibrium and facilitate a transition to a payoff dominated equilibrium. This is not a concern for pure coordination games, in which all players coordinate on the same action or they obtain zero payoffs.

Proposition 5. *When Γ is a pure coordination game with more than a single strict Nash equilibrium:*

⁴ The reasoning in this example extends to more complex settings of multiplayer contracting.
⁵ Movements from the football convention to the opera convention require random errors, but movements in the other direction do not: $r(OOO, FFF) < (0, 0, 1) < r(FFF, OOO)$. That the football convention is CSS then follows by Proposition 3.

	L	R
a	4, 4, 4	2, 0, 6
b	0, 6, 2	5, 5, 0

A

	L	R
a	6, 2, 0	0, 1, 5
b	5, 0, 1	3, 3, 3

B

Fig. 3. A three player strategic game, in which player 3 chooses A or B.

- (i) Where one equilibrium is Pareto superior to all other equilibria, the superior equilibrium is selected by CSS.
(ii) When an equilibrium is Pareto dominated by any other equilibrium, it is never selected by CSS.

Consider a bargaining frontier given by a concave decreasing function $f(z)$, $0 \leq z \leq \hat{z}$, $f(\hat{z}) = 0$, $\hat{z} > 0$. Let us consider a two-player pure coordination game Γ_C in which the available contracts on the frontier are those in which players 1 and 2 receive payoffs $(\kappa\delta, f(\kappa\delta))$ for $\kappa \in \mathbb{N}$, $0 \leq \kappa\delta \leq \hat{z}$, δ assumed small and such that \hat{z} is an integer multiple of δ ; and all other contracts are Pareto dominated by at least one contract on the frontier.

Proposition 6. Define the Nash bargaining solution as z^* such that $zf(z)$ is maximized. When $\Gamma = \Gamma_C$ then CSS selects a contract $(z, f(z))$ such that $|z - z^*| < 2\delta$.

Standard stochastic stability selects the Kalai–Smorodinsky bargaining solution in such games (Young, 1998a). Naidu et al. (2010) attain the Nash bargaining solution by assuming that switches between conventions are always driven by the errors of the player who gains from the switch. Proposition 6 achieves this effect endogenously: players are always willing to switch conventions to one in which they do better, but in order for a player to be encouraged to participate in a coalitional deviation to a convention in which he does worse requires his status quo payoff to be reduced by the random errors of the player who gains from the switch. It should be noted that Propositions 4, 5 and 6 hold even if the probability of coalitional behavior is not taken to become vanishingly small.

The question naturally arises as to how far we can extend our results regarding the preference of CSS for efficiency.

4.5. Efficiency with > two players

In the battle of the sex ratio the selected equilibrium payoff vector is $(2, 2, 6)$. This is an element of the core of the game, defined as in Aumann and Peleg (1960), whether the core for games with non-transferable utility is described using the concept of α -efficiency or β -efficiency.⁶ It is a well-known property of the core that its elements are efficient outcomes of the underlying game. Can we establish any kind of inclusion relation between CSS and the core? The answer is no, we cannot. Even when CSS selects a unique equilibrium we cannot guarantee that this equilibrium is contained in the core of the game. This result is in contrast to Konishi and Ray (2003) where a farsighted dynamic process always selects payoffs in the core of the game when a unique limit of the process exists. Serrano and Volij (2008) demonstrate that stochastic stability does not necessarily select equilibria in the core.⁷ Here I give an example of a game with a nonempty singleton core and a singleton CSS set which are not the same.

In the game in Fig. 3 the unique element of the core is aLA with payoffs $(4, 4, 4)$. CSS chooses bRB – the only inefficient pure strategy combination possible! The reason for this is that it requires a three-player coalition to move from bRB to aLA , whereas a two-player coalition will deviate from aLA to bRA from where a best response of player 3 is to move to bRB (all usual provisos about sample sizes in the adaptive process apply). aLA dominates all other strict Nash equilibria and is Pareto efficient: it therefore follows that Proposition 4 cannot be extended to games with more than two players.

4.6. Coalition proofness

Nor is there an inclusion relation between CSS outcomes and Coalition Proof outcomes (Bernheim et al., 1987). The game in Fig. 3 has bRB as its unique CSS outcome and aLA as its unique Coalition Proof outcome.

aLA is coalition proof as all coalitional deviations lead to further deviations by subsets of the deviating players.⁸ bRB is not coalition proof as the players can jointly deviate to aLA which is itself coalition proof. However, although a deviation from aLA to bRA is disturbed by further deviations, it still allows the possibility of a transition to bRB by single player best responses. Thus, coalitions that are not viable deviations in the Coalition Proof Nash equilibrium concept can change outcomes in the CSS concept if they open up opportunities to enter the basin of attraction of another equilibrium.

The myopia of the players in this paper means that coalitional deviations are feasible even if they are not themselves robust to further coalitional deviations. Interpreted from the point of view of a population model, players selected from the

⁶ α -efficiency guarantees coalitions payoffs at least as high as their maximin payoffs, β -efficiency guarantees coalitions payoffs at least as high as their minimax payoffs. The core under α -efficiency is the set $\{(2, 2, 6), (3, 3, 4)\}$; the core under β -efficiency is $\{(2, 2, 6)\}$.

⁷ This is clearly true for games with non-transferable utility. There is however more reason to suspect that CSS and the core might be related, both being defined using coalitional concepts.

⁸ Deviation to bRA leads to deviation to bLA .

population to play the game in question do not worry about the potential effect of their actions on future players in the same position in the game. As a consequence, CSS deals with a larger class of coalitional deviations than does Coalition Proof Nash equilibrium, since the latter concept does not require robustness to coalitional deviations which can themselves be destabilized by further deviations by subsets of the set of deviating players.

It should be noted that this is the first example in the paper for which switching the ordering of different sizes of coalitional deviations makes a difference to equilibrium selection. Specifically, if three-player deviations were more likely than two-player deviations then *aLA* would be selected.

4.7. Supermodular games

A class of games is now given for which CSS states are efficient and coincide with coalition proof Nash equilibria. If the strategy sets X_i are ordered sets and the payoff functions exhibit *increasing differences*:

$$\tilde{x} \geq x \implies \pi_j(\tilde{x}_j, \tilde{x}_{-j}) - \pi_j(x_j, \tilde{x}_{-j}) \geq \pi_j(\tilde{x}_j, x_{-j}) - \pi_j(x_j, x_{-j})$$

then Γ is *supermodular* and has a largest and a smallest Nash equilibrium in pure strategies. *Positive spillovers* exist if $\pi_j(x_j, x_{-j})$ is increasing in x_{-j} , *negative spillovers* exist if $\pi_j(x_j, x_{-j})$ is decreasing in x_{-j} . With positive (negative) spillovers the largest (smallest) Nash equilibrium is Pareto efficient within the set of Nash equilibria.

Proposition 7. *If Γ is supermodular with positive (negative) spillovers and the largest (smallest) Nash equilibrium is strict, then CSS uniquely selects the largest (smallest) Nash equilibrium.*

These equilibria are also the unique coalition proof Nash equilibria (Milgrom and Roberts, 1994; Moreno and Wooders, 1996). The same result holds if Γ is *quasisupermodular* (Milgrom and Shannon, 1994). It is possible to state a broader result that implies Proposition 7.

Proposition 8. *If Γ exhibits positive (negative) spillovers, has a strict Nash equilibrium x^* that Pareto dominates all other pure Nash equilibria, and has increasing differences for $\tilde{x}_{-j}, x_{-j} \geq x_{-j}^*$ ($\tilde{x}_{-j}, x_{-j} \leq x_{-j}^*$), then x^* is selected by CSS.*

The proof of Propositions 7 and 8 use the fact that even without the probability of coalitional behavior being taken to zero, in the presence of positive (negative) spillovers, any coalitional response to samples of strategy profiles weakly higher (lower) than x^* will itself be weakly higher (lower) than x^* . That is, efficiency bounds are obtained even before coalitional behavior becomes small.

In the following example there is a Pareto dominant Nash equilibrium which is not selected as there are neither unambiguously positive nor negative spillovers. It highlights a counterintuitive implication of coalitional behavior in a classic industrial organization example.

4.8. Example: Bertrand game with convex costs

Consider a model of Bertrand competition with a single undifferentiated good. There are n firms producing the good. Firms simultaneously set prices $p \in \mathbb{N}$ and⁹ the lowest price seller gets all of the demand with demand being shared equally if there are two or more firms with equally low prices. Total demand for the good is D and the cost to a firm of producing d units of the good is given by $ad + bd^2$ where a and b are strictly positive constants. For simplicity we assume that D is not affected by the price p .

Clearly, all strict Nash equilibria of this game involve every firm charging the same price. Some of these Nash equilibria are vulnerable to subsets of players dropping their prices by 1 and sharing the market between themselves; in fact there is a threshold price p' above which it is worthwhile for k players to do just that. This threshold is given by

$$\frac{D}{n}p' - a\frac{D}{n} - b\left(\frac{D}{n}\right)^2 = \frac{D}{k}(p' - 1) - a\frac{D}{k} - b\left(\frac{D}{k}\right)^2$$

which gives

$$p' = a + b\frac{D}{n} + b\frac{D}{k} + \frac{n}{n-k},$$

which is minimized over k at

$$k^* = \frac{n\sqrt{bD}}{\sqrt{n} + \sqrt{bD}},$$

⁹ Note that although here we use discrete prices, the arguments of this section are equally valid for the continuous pricing Bertrand model.

giving

$$p^* = a + 1 + \frac{2bD + 2\sqrt{bDn}}{n}.$$

So if, for example, $n = 10$, $a = 10$, $b = 0.1$, $D = 100$, then price equals average cost at $p = a + b\frac{D}{n} = 11$ and $p'(1) = 22.11$, suggesting that any price between 11 and 22 sustains a Nash equilibrium. However, $k^* = 5$ and $p^* = 15$, so every Nash equilibrium at prices $p \geq 15$ is vulnerable to deviations by coalitions of five players. The set of strategies can be made finite by only looking at prices up to 25 and then this game can be analyzed in an adaptive play setting. CSS then selects a convention as follows: the conventions where $p = 12, 13, 14$ all require ten-player deviations or random errors for the process to leave them, whereas they can be reached by deviations of nine players or fewer from the other conventions. Hence by Proposition 3, the conventions where $p = 12, 13, 14$ are the only candidates for CSS states. The easiest way out of these recurrent classes by ten-player deviations is when the deviations in question involve all firms setting $p = 25$. After this has occurred enough times it is then possible for a nine-player deviation to respond with $p = 24$. This transition out of the recurrent class will be harder the higher the price being charged to begin with, and via these transitions it is possible to reach the $p = 14$ convention, therefore $r(14, \cdot) > r(12, 14), r(13, 14)$ and by Proposition 1 it must be that the convention where all firms set price $p = 14$ is the unique CSS state.

As the number of firms in the market grows $p'(1) \rightarrow a + bD + 1$. In the example above this means $p'(1) \rightarrow 21$ which suggests that the set of Nash equilibria is not significantly reduced in size. However, $p^* \rightarrow a + 1$: as the number of firms becomes large the only Nash equilibrium robust to coalitional deviations of $(n - 1)$ players or fewer and hence the only CSS equilibrium is that where $p = a + 1$. Competitive outcomes occur because of the presence of collusive behavior by firms. Moreover, although $k^* \rightarrow \infty$, $\frac{k^*}{n} \rightarrow 0$: the number of firms involved in the coalitional deviation which persists at the lowest prices increases, but that number as a share of the total number of firms in the market decreases.

Finally, we note that in this example, as in the game in Fig. 3, switching the ordering of different sizes of coalitional deviations makes a difference to equilibrium selection. Specifically, if the ordering is reversed so that larger coalitions are more likely to form than smaller ones, then the Nash equilibrium with the highest price will be selected.

4.9. Structure of coalitional deviations

In my definition of CSS I assume that all possible coalitions have the chance to deviate. This can easily be altered to model situations where certain players are not expected to cooperate with one another. An example of this might be a game with a set of buyers and a set of sellers, where sets of sellers can make coalitional deviations (modeling collusive behavior) but no other set of players can. We can in fact define any hierarchy of subsets of players ordered by the likelihood of the occurrence of coalitional behavior in them:

$$\xi_1, \xi_2, \dots, \xi_M, \quad M \in \mathbb{N}, \quad \xi_i \subset \mathcal{P} \forall i.$$

Naturally, some coalition structures can be considered more reasonable than others. Suggestions have been made in the cooperative game theory literature that subsets of coalitions which are allowed to deviate should also be allowed to deviate^{10,11} or alternatively that the union of coalitions which are allowed to deviate and have a nonempty intersection should also be allowed to deviate.¹² In the first case it is argued that it is possible for a subset of a set of players who meet to discuss strategy to meet without the others present. The second case is predicated on the argument that players who are members of the intersection between two coalitions can serve as intermediaries to bring the interests of the two coalitions together. It is also possible to represent feasible coalition structures using graph theory, for instance by assuming that any coalitional activity is undertaken by connected subgraphs of a graph representing a wider social network.¹³

The author is aware that the model of the current paper could be considered more realistic if payoff considerations were allowed to affect the probability of a coalition's formation. Furthermore, the lexicographic ordering of the coalition sizes could be relaxed: one possible avenue of interest would be to obtain bounds on predictions when the relative coalition formation probabilities are restricted to lie within certain plausible bounds, what constitutes a plausible bound being an empirical question. The relevant history, anthropology and sociology literature¹⁴ should be a fruitful source of inspiration in answering this question and in the adaptation of the model to explain various real world phenomena.

¹⁰ Of course the argument is not phrased in this manner in cooperative papers. Instead the structure of allowable deviations in cooperative games is described by the set of characteristic function inequalities that need to be satisfied in order for a game to be counted as part of the core of the game or to satisfy another cooperative solution concept such as the nucleolus (Schmeidler, 1969).

¹¹ Algaba et al. (2004).

¹² Algaba et al. (2001).

¹³ Myerson (1977), Jackson and Wolinsky (1996), Jackson (2005), Kets et al. (2011).

¹⁴ For example Olson (1971), Chamberlin (1974), Poteete and Ostrom (2004), Mathew and Boyd (2011).

5. Conclusion

This paper has presented the use of coalitional stochastic stability as a method of equilibrium selection, and argued that it should be preferred to random error based stochastic stability wherever coalitional behavior is feasible. We have demonstrated that the ideas underlying CSS are as intuitive as those underlying standard stochastic stability and shown how CSS states can be found. CSS is a way of incorporating coalitional considerations into equilibrium selection and evolutionary game theory and therefore lies at the crossroads of several different strands of literature. We have shown how despite a strong preference for efficiency in the description of CSS, efficiency will not always be attained and that sometimes social movements to other Nash equilibria (such as the French Revolution or the move from *bRB* to *aLA* in Fig. 3) will quickly collapse due to further deviations. This does not mean that such changes will not happen. They will happen, and the resulting instability may take you somewhere new.

In the writing of this paper it was discovered that coalitional perturbations of the type described can lead to faster convergence to long run equilibria than is the case with random error models of stochastic stability. This is due to coordination over: (i) the time when individuals deviate, and (ii) the strategies they choose when they do deviate. Perhaps more surprisingly, in local interaction models coalitional behavior can also lead to a conservative effect, increasing convergence times. Combining the reforming and conservative aspects of coalitional play leads to a ‘tipping point’ effect in which the speed of adoption of a new behavior can be very sensitive to game parameters. Due to space constraints, these results are not included in the current paper.

Coalitional behavior is something that can be observed in many noncooperative games and in discussions of Nash and other equilibrium concepts it is often the elephant in the room: the question being how to incorporate the realism of coalitional behavior without discarding the precision of equilibrium predictions. This paper has given one way of overcoming this problem. It would be interesting to see further analysis of games with coalitions and a multiplicity of equilibria using the tools described in this paper.

Acknowledgments

I would like to thank Robert Evans, Jean-Paul Carvalho, Sanjeev Goyal, Christopher Harris, Matthew Jackson, Michihiro Kandori, David Myatt and Larry Samuelson for their kind advice and comments at various stages of this project. I also thank two anonymous referees and the editorial staff for their excellent criticism.

Appendix A. Proofs

Proof of Proposition 1. For given $b = (b_2, \dots, b_{n+1})$, $h_1, h_2 \in X^m$, define:

$$\tilde{r}(h_1, h_2) = \min \left\{ \hat{r} \in \mathbb{R}_+ : \exists t \in \mathbb{N} \text{ s.t. } \lim_{\epsilon \rightarrow 0} \frac{P_\epsilon^t(h_1, h_2)}{\epsilon^{\hat{r}}} > 0 \right\}.$$

Write stochastic potential of state h under resistances \tilde{r} as $\tilde{S}P(h)$ and under resistances r as $SP(h)$. We know from Young (1993) that:

$$\left(\lim_{\epsilon \rightarrow 0} \mu_\epsilon(h^*) > 0 \right) \iff \left(h^* \in \arg \min_{h \in X^m} \tilde{S}P(h) \right)$$

and that $\lim_{\epsilon \rightarrow 0} \mu_\epsilon(\cdot)$ is a stationary distribution of P_0 . Then setting:

$$\hat{\beta} = \max_{h_1, h_2 \in X^m} \sum_{i=2}^{n+1} |r_i(h_1, h_2)|$$

we have, with “ \cdot ” denoting a dot product:

$$\forall \beta > \hat{\beta}, \forall h_1, h_2 \in X^m: \tilde{r}(h_1, h_2) = r(h_1, h_2) \cdot b.$$

Set $\hat{\beta} = |X^m| \hat{\beta}$. Then $\forall \beta > \hat{\beta}$:

$$\forall h \in X^m: \tilde{S}P(h) = SP(h) \cdot b$$

and

$$\left(h^* \in \arg \min_{h \in X^m} \tilde{S}P(h) \right) \iff \left(h^* \in \arg \min_{h \in X^m} SP(h) \right). \quad \square$$

Proof of Proposition 2. Given that CSS selects amongst states which have positive probability under stationary distributions of P_0 and that we know from Young (1993) that stationary distributions of P_0 have all probability weight on convention(s) under the conditions given in Proposition 2, the proof is complete. \square

Proof of Proposition 3. In Ellison (2000) the radius $R(\cdot)$ and coradius $CR(\cdot)$ of C_1 are defined as:

$$R(C_1) = \min_{\substack{h_1 \in C_1 \\ h_2 \in C_2}} \tilde{r}(h_1, h_2); \quad CR(C_1) = \max_{h_2 \in C_2} \min_{h_1 \in C_1} \tilde{r}(h_2, h_1).$$

As

$$(\forall h_1 \in C_1, \forall h_2 \in C_2: r_k(h_1, h_2) > 0) \implies R(C_1) > b_k$$

and

$$(\forall h_2 \in C_2: \exists h_1 \in C_1: \forall i \geq k: r_i(h_2, h_1) = 0) \implies \exists \alpha \in \mathbb{N}: CR(C_1) < \alpha b_{k-1}$$

then

$$\beta > \alpha \implies R(C_1) > CR(C_1) \implies \left(\lim_{\epsilon \rightarrow 0} \mu_\epsilon(h) > 0 \implies h \in C_1 \right),$$

where the last implication follows from Theorem 1 of Ellison (2000). \square

Proof of Proposition 4. As the conditions of Proposition 2 are satisfied, all CSS states are conventions. Let x^* be the Pareto dominant strict Nash equilibrium, $h^* = (x^*)^m$, $C_1 = \{h^*\}$. Let C_2 contain all other conventions. Then random errors are required to reach C_2 from h^* but h^* can be reached from C_2 without random errors:

$$\forall h_2 \in C_2: r(h^*, h_2) \geq (0, 1), r(h_2, h^*) < (0, 1)$$

so C_1, C_2 satisfy the conditions of Proposition 3 and the unique CSS state must be h^* . \square

Proof of Proposition 5. Pure coordination games are weakly acyclic, so (i) follows directly from the proof of (ii). To prove (ii) let C_1 be the set of Pareto undominated equilibria and C_2 be all other equilibria. Then:

$$\forall h_1 \in C_1, \forall h_2 \in C_2: r(h_1, h_2) \geq (0, \dots, 0, 1)$$

and

$$\forall h_2 \in C_2: \exists h_1 \in C_1: r(h_2, h_1) < (0, \dots, 0, 1)$$

so C_1, C_2 satisfy the conditions of Proposition 3 and the unique CSS state must be contained in C_1 . \square

Proof of Proposition 6. By Proposition 5 any CSS state must be on the bargaining frontier. From any convention on the frontier, random errors are required to move to another convention. Consider a convention where players 1 and 2 receive z and $f(z)$ respectively. One way to leave the convention is for a player to make random errors until the other player best responds in a way that differs from the current convention. If player 2 makes errors and demands nothing then the easiest way to induce player 1 to best respond differently is attained:

$$r_{n+1} \hat{z} \geq (1 - r_{n+1})z \implies r_{n+1} \geq \frac{z}{z + \hat{z}} =: \xi_1(z).$$

When player 1 errs and player 2 responds we similarly obtain:

$$r_{n+1} \geq \frac{f(z)}{f(z) + f(0)} =: \xi_2(z).$$

In the absence of coalitional deviations, equating these two expressions shows that the hardest conventions from which to escape approximate the Kalai–Smorodinsky bargaining solution. However, the possibility of coalitional deviations introduces new ways to escape from conventions, even though the conventions in question are Pareto efficient. The occurrence of random errors can make a coalitional deviation profitable. Analyze a situation where random errors lead to a coalitional deviation $(\tilde{z}, f(\tilde{z}))$, $\tilde{z} < z$. Player 2 will be willing to make these deviations even without sampling any random errors by player 1. We calculate how many random errors player 2 must make to induce player 1 to accept this deviation:

$$\tilde{z} \geq (1 - r_{n+1})z \implies r_{n+1} \geq \frac{z - \tilde{z}}{z},$$

which is minimized when $\tilde{z} = z - \delta$, giving

$$r_{n+1} \geq \frac{\delta}{z} =: \xi_3(z).$$

Similarly for moves to $z + \delta$:

$$r_{n+1} \geq \frac{f(z) - f(z + \delta)}{f(z)} =: \xi_4(z).$$

For small δ : $\xi_3, \xi_4 < \xi_1, \xi_2$ except for values of z close to 0 or \hat{z} . Let $\xi(z) = \min_i \{\xi_i(z)\}$. We construct a spanning tree where every convention is exited by a path of least resistance. All Pareto dominated conventions exit to conventions on the frontier. Conventions where $\xi = \xi_1$ or $\xi = \xi_2$ can reach any other convention by a path with resistance ξ , so we link them by directed edges to conventions for which $\xi = \xi_3$ or $\xi = \xi_4$. From these conventions we add directed edges with resistance ξ to conventions where player 1's payoff differs by δ . ξ_2, ξ_3 decrease in z ; ξ_1, ξ_4 increase in z . The root of a minimum cost spanning tree and thus the CSS state is the convention for which ξ is largest and will be within δ of where $\xi_3(z) = \xi_4(z)$. Also, using the concavity of $f(\cdot)$:

$$\begin{aligned} \xi_3(z) = \xi_4(z) &\implies (z - \delta)f(z) - zf(z + \delta) = 0 \\ &\implies (z - \delta)f(z) - z(f(z) + \delta f'(z)) < 0 \\ &\implies f(z) + zf'(z) > 0 \end{aligned} \tag{A.1}$$

and

$$\begin{aligned} (z - \delta)f(z) - zf(z + \delta) = 0 &\implies (z - \delta)(f(z + \delta) - \delta f'(z + \delta)) - zf(z + \delta) > 0 \\ &\implies f(z + \delta) + (z - \delta)f'(z + \delta) < 0 \\ &\implies f(z + \delta) + (z + \delta)f'(z + \delta) < 0 \end{aligned} \tag{A.2}$$

and together (A.1) and (A.2) imply that the point where $\xi_3(z) = \xi_4(z)$ is within δ of where $f(z) + zf'(z) = 0$. So the CSS convention is within 2δ of where $f(z) + zf'(z) = 0$. \square

Proof of Propositions 7 and 8. Let Γ be a game with positive spillovers. Let x^* be the Pareto dominant strict Nash equilibrium, $h^* = (x^*)^m$, $C_1 = \{h^*\}$. Let C_2 contain all other conventions. Then:

$$\forall h_2 \in C_2: r_{n+1}(h^*, h_2) \geq 1$$

as positive spillovers and increasing differences imply that when samples solely contain strategy profiles $\geq x^*$, coalitional deviations will only choose strategies $\geq x^*$. To see this, let \tilde{x}, \bar{x} be strategy profiles, $\bar{x} \geq x^*$. Let $\hat{x} = \tilde{x} \vee x^*$ where \vee denotes the componentwise maximum. Then for any $Q \subseteq N$:

$$\pi_i(\hat{x}_Q, \bar{x}_{-Q}) \geq \pi_i(\tilde{x}_i, \hat{x}_{Q \setminus \{i\}}, \bar{x}_{-Q}) \geq \pi_i(\tilde{x}_Q, \bar{x}_{-Q})$$

where the first inequality follows from increasing differences and the second from positive spillovers. As x^* is a strict Nash equilibrium, the first inequality is strict if $\hat{x}_i \neq \tilde{x}_i$. Also note that from any history of strategy profiles $\geq x^*$ the process P_0 converges to x^* . From any dominated equilibrium, h^* can be reached by n -player coalitional deviations without any random errors:

$$\forall h_2 \in C_2: r_{n+1}(h_2, h^*) = 0$$

so C_1, C_2 satisfy the conditions of Proposition 3 and the unique CSS state must be h^* . A similar argument applies for negative spillovers. \square

References

Algaba, E., Bilbao, J., Van Den Brink, R., 2004. Cooperative games on antimatroids. *Discrete Math.* 282, 1–15.
 Algaba, E., Bilbao, J.M., Borm, P., López, J.J., 2001. The Myerson value for union stable structures. *Math. Methods Operations Res.* 54, 359–371.
 Ambrus, A., 2009. Theories of coalitional rationality. *J. Econ. Theory* 144, 676–695.
 Aumann, R., 1959. Acceptable points in general cooperative n -person games. In: Tucker, A.W., Luce, R.D. (Eds.), *Contributions to the Theory of Games*, vol. IV. Princeton University Press, pp. 287–324.
 Aumann, R., Peleg, B., 1960. Von Neumann–Morgenstern solutions to cooperative games without side payments. *Bull. Amer. Math. Soc.* 66.
 Bergin, J., Lipman, B.L., 1996. Evolution with state-dependent mutations. *Econometrica* 64, 943–956.
 Bernheim, B.D., Peleg, B., Whinston, M.D., 1987. Coalition-proof Nash equilibria I. *Concepts. J. Econ. Theory* 42, 1–12.
 Chamberlin, J., 1974. Provision of collective goods as a function of group size. *Amer. Political Sci. Rev.* 68, 707–716.
 Ellison, G., 2000. Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. *Rev. Econ. Stud.* 67, 17–45.
 Foster, D., Young, H.P., 1990. Stochastic evolutionary game dynamics. *Theoret. Popul. Biol.* 38, 219–232.
 Hwang, S.H., 2011. Larger groups may alleviate collective action problems. Santa Fe Institute Working Paper Series, 09-08-035.
 Jackson, M.O., 2005. Allocation rules for network games. *Games Econ. Behav.* 51, 128–154.
 Jackson, M.O., Wolinsky, A., 1996. A strategic model of social and economic networks. *J. Econ. Theory* 71, 44–74.
 Kandori, M., Mailath, G.J., Rob, R., 1993. Learning, mutation, and long run equilibria in games. *Econometrica* 61, 29–56.
 Kets, W., Iyengar, G., Sethi, R., Bowles, S., 2011. Inequality and network structure. *Games Econ. Behav.* 73, 215–226.

- Konishi, H., Ray, D., 2003. Coalition formation as a dynamic process. *J. Econ. Theory* 110, 1–41.
- Luo, X., Yang, C.C., 2009. Bayesian coalitional rationalizability. *J. Econ. Theory* 144, 248–263.
- Mathew, S., Boyd, R., 2011. Punishment sustains large-scale cooperation in prestate warfare. *Proc. Natl. Acad. Sci.* 108, 11375–11380.
- Milgrom, P., Roberts, J., 1994. Strongly coalition-proof equilibria in games with strategic complementarities. Stanford University, Department of Economics, Working Papers 95002.
- Milgrom, P., Shannon, C., 1994. Monotone comparative statics. *Econometrica* 62, 157–180.
- Moreno, D., Wooders, J., 1996. Coalition-proof equilibrium. *Games Econ. Behav.* 17, 80–112.
- Myerson, R.B., 1977. Graphs and cooperation in games. *Math. Operations Res.* 2, 225–229.
- Naidu, S., Hwang, S.H., Bowles, S., 2010. Evolutionary bargaining with intentional idiosyncratic play. *Econ. Letters* 109, 31–33.
- Olson, M., 1971. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Harvard Economic Studies, revised edition, vol. 124. Harvard University Press.
- Peleg, B., Sudholter, P., 2003. *Introduction to the Theory of Cooperative Games*. Kluwer Academic, Boston, ISBN 1-4020-7410-7. 378 pp.
- Poteete, A.R., Ostrom, E., 2004. Heterogeneity, group size and collective action: the role of institutions in forest management. *Devel. Change* 35, 435–461.
- Schmeidler, D., 1969. The nucleolus of a characteristic function game. *SIAM J. Appl. Math.* 17.
- Serrano, R., Volij, O., 2008. Mistakes in cooperation: the stochastic stability of Edgewood's recontracting. *Econ. J.* 118, 1719–1741.
- Smith, J.M., Price, G.R., 1973. The logic of animal conflict. *Nature*, 15–18.
- van Damme, E., Weibull, J.W., 2002. Evolution in games with endogenous mistake probabilities. *J. Econ. Theory* 106, 296–315.
- Young, H.P., 1993. The evolution of conventions. *Econometrica* 61, 57–84.
- Young, H.P., 1998a. Conventional contracts. *Rev. Econ. Stud.* 65, 773–792.
- Young, H.P., 1998b. *Individual Strategy and Social Structure*. Princeton University Press.