



Quasi-Conventions

Brian Skyrms¹

Received: 4 October 2021 / Accepted: 17 February 2023 / Published online: 3 March 2023
© The Author(s) 2023

Abstract

I consider a generalization of Vanderschraaf's correlated conventions to Quasi-Conventions, using the concept of coarse correlated equilibria. I discuss the possibility of improved payoffs and the question of learnability by simple uncoupled learning dynamics. Laboratory experiments are surveyed. The generalization introduces strains of commitment, which can be seen from different points of view. I conclude that the strains of commitment preclude using the generalization as a stand-alone definition of convention, but that in certain settings Quasi-Conventions can be important modules within larger true conventions.

Keywords Convention · Correlated equilibrium · Coarse correlated equilibrium · Commitment · Resolute choice · Utilitarianism

Lewis (1969) defined a social convention as a behavioral regularity that can be viewed as a certain kind of equilibrium in a game, a game that models the social interaction in which that convention is operative. A convention is conventional in that there are multiple equilibria in the game that are potential conventions. A convention can be maintained because it is an equilibrium, in Lewis a kind of strict Nash equilibrium.

Vanderschraaf (1995, 1998) developed an extension of Lewis' treatment that moves from Nash equilibrium to Aumann's (1974, 1987) *correlated equilibrium*. See Vanderschraaf (1998) for a comprehensive and rigorous development of his theory and comparison with a careful reconstruction of Lewis' informal theory. He uses his theory of convention in his (2019) book on strategic justice. In correlated equilibria players may condition their behavior on some external random feature of their environment. Vanderschraaf's theory of correlated convention retains the virtues of Lewis' account, and adds to them. Conventionality and stability are preserved. Vanderschraaf allows for more conventions and covers natural examples that are excluded by Lewis' account. Some of his correlated conventions are more efficient than any Lewis convention in

✉ Brian Skyrms
bskyrms@uci.edu

¹ Department of Logic and Philosophy of Science, University of California, Irvine, Irvine 92697-5100, USA

that they lead to greater payoffs for all parties. Recent work in computer science shows that they are, in a sense, *theoretically* easier to learn. Laboratory experiments demonstrate that people do, in fact learn them.

Here, I would like to consider the prospect of moving a little further down this path. That is, to consider the move from correlated equilibrium to an even more general equilibrium concept, *coarse correlated equilibrium* (Moulin & Vial, 1978), that has recently been a focus of interest in learning dynamics in algorithmic game theory. As in correlated equilibria, players can adopt strategies conditional on external random features of the world. You can visualize the external random feature of the world on which players key their acts as a lottery run by a mediator. The mediator runs the lottery, observes the outcome, and conditional on the outcome makes a recommendation to each player. The players know the structure of the lottery; they know the joint probability of recommendations. Players can either choose either to commit to the lottery or to just take an action in the game. The lottery is a *coarse correlated equilibrium* if when every player *commits* to following the recommendations of the the lottery, each is making a best response to the others. In other words, no one would have had an expected payoff higher than that given by the lottery if she had chosen otherwise while all others chose the lottery. (This is more permissive than a correlated equilibrium, which requires that once she gets the recommendation, honoring it is still a best response.)

The purpose of this paper is to review this concept and to ask whether the move from Vanderschraaf's definition of convention in terms of *correlated equilibrium* to an analogous definition in terms of *coarse correlated equilibrium* gives us an even better characterization of a convention. I will conclude that it does not, but that nevertheless it can have an important subsidiary role to play in the analysis of conventions. For these reasons, I call the resulting analogue of a correlated convention, a *Quasi-Convention*.

All correlated equilibria are coarse correlated equilibria but coarse correlation allows for even more possibilities. Some of these possibilities are attractive, but Quasi-Conventions lose a cardinal virtue of the Lewis-Vanderschraaf conventions. They are not stable. They come with *strains of commitment*, akin to those of Utilitarianism. On the other hand, they may, in the best case, have even better payoffs for all players than Correlated Conventions. On the issue of learnability, results are mixed. They are, in a sense, *theoretically* even easier to learn. But existing laboratory experiments fail to demonstrate that people do learn them. *Quasi-Conventions* thus fall short of being robust extensions of the concept of convention, but they are not without interest.

Section 1 will review the relevant equilibrium concepts. Section 2 gives the definitions of convention in Lewis and Vanderschraaf and extension to *Quasi-Conventions*. Section 3 discusses the strains of commitment, how they figure in Rawls and Binmore's objections to Utilitarianism, how they apply to Quasi-Conventions, and how McClenen's theory of resolute choice discounts them. Section 4 compares learnability of the different kinds of equilibria using various learning dynamics. Section 5 reviews the empirical evidence of learnability from laboratory experiments. Section 6 sums up and discusses the status of *Quasi-Conventions* in light of the foregoing.

1 Equilibria

We have a hierarchy of equilibrium concepts; each class is a subset of those below it in this list:

1. Pure strategy Nash equilibria
2. Mixed strategy Nash equilibria
3. Correlated equilibria
4. Coarse correlated equilibria

1.1 Pure strategy Nash and strict Nash equilibria

Consider normal form games with a finite number of players, each with a finite number of possible acts. The vector of payoffs for the players is a function of the vector of acts for each player, the act profile. The set of players, the set of acts available for each player and the payoff function jointly define a normal-form game. An act profile is a Nash equilibrium if each player is playing a best response to the joint play of the others. In other words, no player could have improved her payoff by changing her act while all others kept theirs the same. This is a (weak) stability concept. No player could gain from such a hypothetical unilateral deviation, but one might be just as well off. If, in addition, any unilateral deviation would have made the deviating player’s payoff strictly worse, it is a Strict Nash Equilibrium. Strict Nash equilibrium gives us a strong sense of stability.

Here are two examples, (which we will revisit):

A version of Chicken (Aumann, 1974):

	Aggressive (Hawk)	Chicken (Dove)
Aggressive (Hawk)	0, 0	7, 2
Chicken (Dove)	2, 7	6, 6

A version of Rock-Scissors-Paper:

	Rock	Scissors	Paper
Rock	− 9, − 9	1, − 1	− 1, 1
Scissors	− 1, 1	− 9, − 9	1, − 1
Paper	1, − 1	− 1, 1	− 9, − 9

Chicken has two strict Nash equilibria in pure strategies, $\langle \text{Hawk}, \text{Dove} \rangle$ and $\langle \text{Dove}, \text{Hawk} \rangle$. Rock-Scissors-Paper has none.

1.2 Mixed strategy Nash equilibria

Now we allow each player to choose a probability distribution over her possible acts. This is a mixed strategy. In computing the payoff vector resulting from a vector of mixed strategies, we form a joint probability distribution *assuming that the act probabilities are independent*. Payoffs are taken to be expected payoffs according to this joint probability distribution. One might picture players choosing settings on independent randomizing devices, and then letting the chips fall where they may.

Nash and Strict Nash equilibria in mixed strategies are defined as before using these payoffs. Nash showed that Nash equilibria always exist in this case, when there are finite numbers of players and strategies. But there are no strict Nash equilibria in mixed strategies that are not pure strategies, because of the independence assumption.

In Rock-Scissors-Paper there is now a mixed equilibrium where each player chooses each act with equal probability. In Chicken, a new mixed equilibrium appears, where each player plays Chicken with probability $2/3$. It is more equitable than either pure strategy equilibrium in that it gives each player equal chance of getting their favorite outcome. But it carries with it a substantial probability of the worst outcome, because of the independence assumption. The expected payoff is $4/3$.

1.3 Correlated equilibria

We lift the independence assumption. There is a joint probability distribution over player's acts with no restrictions. We can think of the joint probability distribution as a *mediator* making suggestions to players. Players know the joint distribution and calculate the expectation *conditional on the action recommended to them*. Nash and Strict Nash are defined by this conditional expectation. If for every action that will be recommended with positive probability, no unilateral deviation does better in the expectation *conditional on this act*, then the joint probability distribution is a correlated equilibrium.

In Chicken, there is a correlated equilibrium which gives each of the pure strategy Nash equilibria probability $1/2$. This restores equity in that it gives each player equal payoff while avoiding the unfortunate consequences of the mixed Nash equilibrium. Each player gets an average payoff of $4/3$. It is also *strict*, which is not possible with mixed equilibria.

In addition, as Aumann (1974) pointed out, it is possible to do even better. If the joint probability distribution gives $\langle \text{Chicken}, \text{Chicken} \rangle$, $\langle \text{Aggressive}, \text{Chicken} \rangle$ and $\langle \text{Chicken}, \text{Aggressive} \rangle$ all probability $1/3$, we also have a correlated equilibrium. This equilibrium is also *strict*. The expected payoff for each player is 5.

1.4 Coarse correlated equilibria

We now lift the assumption of conditioning on your own act. A joint coarse correlated equilibrium is a joint probability distribution such that for each player, the expected payoff is greater than that which would come from a unilateral deviation to some pure act. *This is just like the definition of mixed Nash equilibrium, without the independence assumption.* One can picture the players turning over their choice of pure act to a joint randomizing device and letting the chips fall where they may.

We could think of the difference between correlated equilibrium and coarse correlated equilibrium respectively as the difference between a mediator and an arbitrator. A mediator makes a private recommendation to each of the players, and each (knowing the mediator’s propensities) decides whether to accept or deviate. An arbitrator takes over the decision. The players decide whether to leave it to the arbitrator, (*knowing her propensities*), or just choose one of the acts.

Consider the following game, from Georgalos et al. (2020):

	X	Y	Z
A	3, 3	1, 1	4, 1
B	1, 4	5, 2	0, 0
C	1, 1	0, 0	2, 5

There is a unique Nash equilibrium at $\langle A, X \rangle$, and there is no additional mixed Nash equilibrium.

The joint probability distribution that puts probability $\frac{1}{2}$ on $\langle B, Y \rangle$ and probability $\frac{1}{2}$ on $\langle C, Z \rangle$ is a coarse correlated equilibrium. It is not a correlated equilibrium. If told to play C, row will prefer to play A; if told to play Y, column would prefer to play X. But if one only knows these joint probabilities both row and column will prefer to stick for an expected payoff of 3.5. It is *strict*. Any unilateral deviation makes the player worse off in expected payoff.

Note that the expected payoff of 3.5 for each player is better than the payoff of 3 at the unique Nash equilibrium.

For an example with a slightly different flavor, consider this game, from Georgalos et al. (2020):

	X	Y	Z
A	3, 2	2, 0	0, 3
B	0, 3	3, 2	2, 0
C	2, 0	0, 3	3, 2

Here there is no Nash equilibrium in pure strategies. The unique mixed Nash equilibrium has each player choosing each of her strategies with equal probability.

Each has an expected payoff of $5/3$.

The joint probability distribution that puts equal probabilities on the diagonal, $1/3$ on $\langle A, X \rangle$, $1/3$ on $\langle B, Y \rangle$, $1/3$ on $\langle C, Z \rangle$, is a coarse correlated equilibrium, but not a correlated equilibrium. It gives expected payoff on 3 to row player and 2 to column player. If either player chose to deviate to a pure strategy, her expected payoff would drop to $3/2$. Thus, this coarse correlated equilibrium is *strict*.

It is not a correlated equilibrium, because of strains of commitment, Conditional on column player being told to play X (for a payoff of 2) she would rather deviate to Z (for a payoff of 3). Likewise, if told to play Y she would rather play X, and if told to play Z she would rather play Y.

Again, this coarse correlated equilibrium improves the expectations of both players over what they could get from mixed Nash.

2 Conventions

For Lewis, conventions are strict Nash equilibria. But Lewis also requires that any unilateral deviation makes *all* players in the game worse off. This extra proviso is not required to explain the stability of convention. In games of pure common interest, the extra proviso would automatically be satisfied. But as Vanderschraaf (1998, 2001) points out, the restriction to such games blocks the extension of the account of convention to other cases of partial conflict where we do have operative social conventions. There the extra proviso may not be satisfied. Vanderschraaf points to the Chicken game of the previous section, and to various examples of conventions given by David Hume. So Vanderschraaf argues for a pared-down account that just uses strict Nash equilibria. For these, a unilateral deviation makes the deviator worse off, nothing said about what it does to others. In either case, there can be no Lewis conventions in mixed strategies because such equilibria cannot be strict. (But see Binmore (2008), who drops the strictness condition in order to allow mixed equilibria to be conventions.) If we are confined to pure strategies, Nash equilibria can fail to exist. All the more so, for Strict Nash equilibria.

Vanderschraaf's account replaces Strict Nash equilibria with Strict Correlated equilibria. There are examples of correlated conventions in Lewis' book, although he does not call them that. One is that of the traffic light. A traffic light is red or green. *Stop on Red, Go on Green* is a correlated equilibrium, provided that whether you hit a red or green is thought of as random. Lewis thinks of these equilibria as just Nash equilibria. A correlated equilibrium is a Nash equilibrium in a different game. A mixed Nash equilibrium can be thought of as a pure strategy Nash equilibrium in a different game, where players choose between an infinite number of randomizing devices to determine their acts in the original game. There is a careful discussion of this in Vanderschraaf (1998). But if we keep the game fixed, these are different concepts that must be distinguished. Mixed Nash equilibria extend pure Nash, and correlated equilibria extend mixed Nash. Strict correlated equilibria extend strict Nash equilibria. In the traffic light game, GO if Green, STOP if red is a Correlated Convention that extends the

conventions available in the original game. It is clear that Vanderschraaf does extend Lewis. *Quasi-conventions* are Strict Nash equilibria in another different game. This is the game where players *commit* to a joint randomizing device, and remove the possibility of choosing a different pure act when the find out the pure act chosen by that device. So Quasi-conventions extend Vanderschraaf conventions.

Are there real-life examples of Quasi-Conventions? Let us return to our much-discussed example, the traffic light. Instead of just being either red or green, there is now a yellow light. The light can be either red, green, or yellow. There are three possible acts: Go right through, proceed slowly and with Caution, and Stop. Here is a plausible payoff matrix:

	Go	Caution	Stop
Go	0, 0	3, 1	7, 2
Caution	1, 3	2.1, 2.1	6, 2
Stop	2, 7	2, 6	4, 4

If the other guy goes, I prefer to stop. If the other guy stops, I prefer to go. If the other guy proceeds with caution, I prefer to go. (If you are worried about individuals approaching from 3 or 4 directions, make it the intersection of two one-way streets.)

Suppose drivers follow the law: stop on red, go on green, caution on yellow. Suppose the traffic light puts 1/3 on Red for driver 1, Green for driver 2, 1/3 for Green for driver 1, Red for driver 2, and 1/6 each for Yellow for 1, Red for 2, and Red for 1, Yellow for 2.

Then the joint correlated strategy puts probability on the 4 bolded cells of the foregoing payoff matrix: 1/3 on <Go, Stop>, 1/3 on <Stop, Go>, and 1/6 on each <Caution, Stop>, and <Stop, Caution>. This is a not a correlated equilibrium, because conditional on seeing the yellow light drivers will prefer to go. It is a coarse correlated equilibrium. Expected payoff of the equilibrium is $(2/6) 2 + (1/6) 2 + (2/6) 7 + (1/6) 6 = 26/6 = 4.34$. Suppose Row player switches to a pure strategy. Row player will meet Go, with probability 1/3, Caution with probability 1/6, and Stop with probability 1/2. Expected payoff of switching to GO is $(2/6) 0 + (1/6) 3 + (3/6) 7 = 24/6 = 4$. Expected payoff of switching to Caution is $(2/6) 1 + 1/6 (2.1) + (3/6) 6 = 3.6834$. Expected utility of switching to Stop is $(2/6) 2 + (1/6) 2 + (3/6) 4 = 16/6 = 2.67$.

There are examples of quasi-conventions in economically interesting settings. Moulin et al. (2014) discuss a class of two-person games with quadratic payoff functions, including examples of duopoly and two-person public goods provision games. For this class of games, correlated equilibria cannot improve on the payoffs of Nash equilibria, but coarse correlated equilibria can.

As an example, they give a public goods provision game where Xena and Yvette contribute x and y , with Xena's payoff being $x + 4y - (x + y)^2$ and Yvette's being the mirror image, $4x + y - (x + y)^2$. If they agree that one of them will contribute 1 and the other 0, and to flip a fair coin to decides who contributes, they are at a coarse

correlated equilibrium. Before the coin flip, each has an expected payoff of 1.5. She can do no better by deviating. But it is not a correlated equilibrium because after learning that she had lost the flip she would not want to contribute more than 0.5. With this arrangement they both do better than any Nash or Correlated equilibrium. Plausibly many arrangements with this structure exist and are honored.

One might argue that they are honored only because they are embedded in bigger games. This may well be true, but this does not deprive them of interest. For example, Coarse correlated equilibria have a relation to a special kind of bigger game, a repeated game with imperfect monitoring.

In a repeated game, the same so-called Stage Game is played over and over, but the payoffs of strategies in the whole repeated game are computed by geometrically discounting the future and summing the infinite series. A player is said to be patient for small discount rates. This introduces the possibility of sustaining efficient equilibria by punishment strategies in the repeated game that would not be equilibria at all in the stage game. For example, cooperation in the repeated Prisoner's Dilemma can be sustained by Tit-for-Tat or Grim Trigger if each player knows what the others did on the previous round, and the players are sufficiently patient. With imperfect monitoring, players do not know what others did in the past, but only get a noisy signal about it. For a fixed game and a fixed discount rate, the quality of the signal can be crucial to what can be sustained in equilibrium.

In a coarse correlated equilibrium, a player can't do better by unilaterally deviating *ex ante*. In an epsilon-coarse correlated equilibrium it is not worth much to deviate. A player can only do epsilon better by deviating. A coarse correlated equilibrium is an epsilon-coarse correlated equilibrium with epsilon equal to zero.

Awaya and Krishna (2019) show that in an n-player repeated game with imperfect monitoring, the payoffs sustainable as Nash equilibria in the repeated game are contained in the payoffs given by the epsilon-coarse correlated equilibria of the stage game, with epsilon being determined by the discount factor and the quality of monitoring. (They also show how adding communication allows better payoffs.) Poor monitoring, with lots of noise, gives small epsilon. For a large class of situations coarse correlation has an important role to play by being embedded in a larger context in which it is not the operative equilibrium concept.

3 Strains of commitment

“Strains of Commitment” is the apt phrase of Rawls (1971). The issue has played an important role in social philosophy ever since the Greeks. Rawls uses it as an argument against Utilitarianism. He argues that a Utilitarian social contract may involve inequalities so great that those at the bottom will not honor the contract. Let us see how this may involve Quasi-conventions.

Lotteries are common in social life. Lotteries are used as a fair device to award indivisible goods. Lotteries can be used to assign pleasant or unpleasant tasks. We might even draw straws to see who gets pushed out of the lifeboat. In typical lotteries winning is better, and losing worse than not playing at all. In many cases, losers have no recourse. But in some cases, losers may be able to opt out by not performing

as task assigned to them. When this is the case, when the lottery is attractive *ex ante* in expectation, but worse than not performing *ex post*, for some participants, then honoring the lottery arrangement is a coarse correlated equilibrium that is not a correlated equilibrium.

Rawls' argument against Utilitarianism is informal; there is no exact model provided. One is provided by Binmore (1994, 1998). Binmore models the problem to be solved by the social contract as that of selecting a Pareto optimal equilibrium in a Nash Demand game. The just choice is to be selected disinterestedly, from behind a veil of ignorance. Following Harsanyi (but not Rawls) the veil incorporates an equal chance of being anyone in society. Harsanyi notes that the expected utility maximizing choice would be the Utilitarian solution to the Nash Demand game. Here Utilitarianism is justified as the outcome of a fair (imaginary) lottery.

So far, the Utilitarian solution can be a Lewis convention. The Utilitarian solution is, after all, a strict Nash equilibrium of the Nash Demand game, as is every point on the Pareto frontier. If a player were to demand less, he would get less. If he were to demand more, demands would be incompatible and no one would get anything. There is no room yet in the model for *ex post* opting out. Binmore provides this by introducing the possibility of any individual at any time calling to renegotiate the social contract. That is going back behind the veil of ignorance, and running the whole process again. In this extended game, a utilitarian solution may not be a Nash equilibrium. If the Utilitarian social contract is inegalitarian, then the option of renegotiating is attractive to those at the bottom. In the extended game Utilitarianism suffers from the strains of commitment while Egalitarianism does not.

Strains of commitment are not a problem for rational individuals if one builds commitment into one's conception of rationality, and McClennen does just this in his (1990) theory of *resolute choice*. The basic idea is that one does not choose individual acts. One chooses a plan of action, and just sticks to that plan. McClennen illustrates with the story of Ulysses and the Sirens. Ulysses and his sailors prefer to hear the Sirens' song and sail straight. The Sirens' song causes a change of preferences, such that one who hears it prefers to sail toward the rocks. Strotz' (1955–1956) *myopic* rational chooser first chooses to listen to the Sirens and then chooses to sail toward the rocks. His *sophisticated* chooser, realizing how the song will change his preferences chooses to either restrict his possibilities for future choice, like Ulysses, or to stop his ears and not hear the song. McClennen's *resolute* chooser plans to hear the song, keep his options open and sail away from the rocks—and simply sticks to this plan. He ignores his changed preferences because they are at variance with his plan. According to McClennen, this is the rational thing to do. Sophisticated choosers are sensitive to the strains of commitment; resolute choosers are impervious.

McClennen is not completely isolated among philosophers. Closely related views are to be found in Gauthier (1986) and Bratman (1987). Resolute choice is discussed as a way out of diachronic incoherence for non-expected utility theory in Buchak (2013). But we must emphasize that there are nuanced differences among these thinkers that are not discussed here.

In a society of resolute choosers, coarse correlated equilibrium is a natural equilibrium concept. Should those who follow McClennen base their social philosophy on Quasi-Conventions? Or are there still strains of commitment? McClennen freely

grants that actual humans may not be resolute choosers. He regards this as one of many common human failings of rationality, and advocates character development aimed at developing a robust capacity for resolute choice.

He would agree that in actual societies, there are strains of commitment. But he would hold that in a society of rational decision-makers there would not be. So, for him Quasi-Conventions would be viable in a rational society.

The question remains whether, in actual societies, Quasi-Conventions will persist. In many social situations, many people do honor their commitments. It may be taken as a matter of honor to honor one's commitments. Suppose, for instance, there is a group where various jobs need to be done. One job is especially onerous, and a member would prefer not to be in the group than have to do this job. The jobs are assigned by some chance device and in expectation, all individuals strongly prefer to be in the group. All commit to doing the onerous job, if unlucky enough to be assigned it. The "lottery" is run and someone gets the short end of the stick. Does he or she just walk away from the group? If the job is not too onerous, many would not. But if it is, the strains may break the commitment. A reader suggests Aesop's fable of belling the cat. The mice agree to conduct a lottery to see who bells the cat. but in the end, the cat remains un-belled. For formal examples, those of Moulin et al. (2014) referenced above are relevant.

There are various ways to think about this. From the point of view of mainstream rational choice, one might insist that there is a larger game in the background that forms the context within which the smaller game is played, and honoring one's commitments in the game modeling the immediate social situation may just be maximizing payoff in the larger game. How onerous a task must be to break to commitment would depend on the particulars of this larger game. But in some cases, the kinds of conventions sustainable in the larger game have important connections to coarse correlated equilibria in the embedded game. The results of Awaya and Krishna (2019) discussed above are a case in point.

A quite different perspective has recently emerged in computer science. The focus of this literature is the optimal design of interacting software agents. These are resolute choosers, not because of any underlying theory of rationality, but just because they do what they are programmed to do. Since, in some situations, coarse correlated equilibria can carry average payoffs greater than Nash or Correlated equilibria, and because of attractive learnability, the study of Coarse Correlated equilibria has become a standard topic in algorithmic game theory (Nissan et al., 2007; Roughgarden, 2016) This literature has no reason to worry about strains of commitment.

4 Learning in theory

Conventions should be learnable. Ideally, realistic conventions should be easily learnable. This is not in general true for Nash equilibria, although it is for certain special cases. In the early days of game theory, Brown (1951) introduced a learning dynamics for games, fictitious play. Players play the game repeatedly and inductively form beliefs about the other's likely moves. The leading idea is that each player takes others' frequencies of past play as a proxy for probabilities of next play. This can be seen as

Bayesian updating using Laplace's rule of succession. (This is spelled out Fudenberg and Levine (1998) and in Vanderschraaf (2001)). Players then choose a best response to their beliefs. The process repeats.

Fictitious play was thought of also as a way of computing Nash equilibria. Robinson (1951) showed that fictitious play converges to a Nash equilibrium for zero-sum games. The sense of convergence needs to be explained. The only Nash equilibria here are mixed, but the players are playing best responses, which are pure strategies. Convergence here means that the relative frequencies of played actions converge to probabilities that constitute a mixed equilibrium. Convergence like this will be important for Correlated and Coarse Correlated equilibria.

Fictitious play has been shown to converge in other classes of games, but Shapley (1964) provided an example where the historical relative frequencies do not converge, but rather keep cycling forever. Shapley problems generalize to a wide class uncoupled learning dynamics (Hart & Mas-Colell, 2003, 2006). For large games, learning Nash equilibria, even when they are learnable, may also be slow. Time to convergence for uncoupled learning dynamics may scale exponentially with the number of players (Hart & Mansour, 2010). Learning Nash equilibria can be hard.

Learning Correlated equilibria is easier. Vanderschraaf (2001) is a pioneering study using a modification of fictitious play. See also Vanderschraaf and Skyrms (2003), where taking turns is investigated as a learnable approximation to correlated equilibria. There are a number of learning dynamics that always converge to Correlated equilibria. Foster and Vohra (1997) use calibrated forecasting rules, Fudenberg and Levine (1998) use smoothed calibrated fictitious play, Hart and Mas-Colell (2000) use regret-based learning. Convergence here means convergence of the relative frequencies of play to the convex set of Correlated Equilibria.

These are all uncoupled dynamics, in which a player needs to know his own payoff function but not those of other players. *Players also only take into account relative frequencies of past plays.* The Hart and Mas-Colell papers referred to in the last paragraph show that no such learning dynamics always converges to Nash. (The italicized proviso is essential to the result. See Shamma and Arslan (2005) and Huttegger (2017, Ch. 3)).

This sort of dynamics may be of special interest for the study of conventions, since it offers an account of how Correlated Conventions might arise spontaneously. Learning Correlated equilibria for large numbers of players is faster. It scales in polynomial time. All in all, consideration of learning dynamics just reinforces the case for extending the definition of convention to Vanderschraaf's Correlated Conventions.

The learning dynamic advantages of Correlated equilibria over Nash equilibria extend, all the more so, to Coarse Correlated equilibria. For both there are simple, uncoupled, dynamics that lead relative frequencies of play profiles to converge to the convex sets, respectively, of Correlated and Coarse Correlated equilibria.

They can both be learned by regret-minimizing learning dynamics, with different notions of regret. *External regret* for a pure strategy on a play is the expected loss from playing the correlated strategy instead of that pure strategy. So, if external regret is zero, one is at a *Coarse Correlated equilibrium*. In the special case that the joint correlated strategy makes the players independent, this is just the definition of a Nash equilibrium in mixed strategies. *Internal regret* on a play is the same, except it takes the expected

loss conditional on the action recommended on that play by the correlation device. Internal regret is zero for any recommendation when at a *Correlated equilibrium*.

There are various regret minimization dynamics guaranteed to drive regret to zero. The simplest one is regret-matching. The idea is to keep track of historical sum of regret for each pure strategy. Then on a given round, play each pure strategy with probability proportional to its accumulated regret. The joint relative frequencies of past play converge to the appropriate convex set of equilibria. For details see Hart and MasColell (2000), Young (2004), Hart (2005), Nisan et al. (2007), and Roughgarden (2016). There is even an algorithm for converting *any* learning dynamics for coarse correlated equilibrium to one for correlated equilibria (Blum & Mansour, 2007). It does extract a cost for speed of learning in large games. Note that here the correlating device is not an external mediator or arbitrator, but rather the history of play. Both kinds of equilibria can be learned in polynomial time, rather than exponential time. The dynamics is a little simpler for Coarse Correlated equilibria, and the learning is a little faster.

However, there are problems with trying to put these positive learning results together with the positive payoff results we saw in Sect. 1. Correlated equilibria may have better payoffs than any Nash equilibrium in a game, but they may also have worse payoffs. Coarse correlated equilibria may have better payoffs than any correlated equilibrium in a game, but they may also have worse payoffs than any correlated equilibrium. In both cases, a bigger net catches better fish and worse fish. So, we should ask whether it is easy for learning to approximate a good correlated equilibrium, rather than just the class of correlated equilibria. And likewise, we should ask whether it is easy for learning to approximate a good coarse correlated equilibrium.

Unfortunately, general results of this nature are negative. Barman and Ligett (2015) show that one cannot in general efficiently compute a correlated equilibrium, or coarse correlated equilibrium with payoffs better than the worst in those classes respectively. They do show that for a certain class of games, there are positive results, and conclude “Developing dynamics that quickly converge to high-welfare CE/CCE in particular classes of games also remains an interesting direction for future work.”

Given these mixed results, we are led to ask whether people can learn to play correlated, or coarse correlated, equilibria, and if so whether they learn to play good and bad ones or only good ones.

5 Learning in the lab

A whole series of laboratory experiments demonstrate the implementability of Correlated Conventions. Different ways of inducing correlation are investigated. Moreno and Wooders (1998) allow pre-play conversation among the players. They use a three-player matching pennies game, in which players 1 and 2 have an incentive for all to match and player 3 has an incentive to prevent this. If players 1 and 2 play as a team they both choose heads or both choose tails. This is then equivalent to two-person matching pennies played between the team and player 3.

Players play the game once, and communicate anonymously through computers. The evidence is that players do not play Nash equilibrium, but do tend to play a correlated equilibrium that is proof against coalitional deviations.

Cason and Sharma (2007) make direct recommendations to the players that constitute a correlated equilibrium with payoffs superior to any in the convex hull of Nash equilibrium payoffs. The game is two-player, a version of Chicken (Hawk-Dove). They did one experiment with two human players, and another experiment where one human player was matched against a robot known to follow recommendations. Two human players did not play Nash equilibrium and demonstrated a tendency toward correlated equilibrium. The correlated equilibrium was more perfectly implemented when the robot player was introduced.

Duffy and Feltovich (2010) address the question raised at the end of the last section. Like Carson and Sharma, they also make direct recommendations and use a version of Chicken. They test the implementability of a “Bad” correlated equilibrium, with payoffs less than those in the convex hull of Nash equilibrium payoffs, as well as a “Good” correlated equilibrium with payoffs greater, and a “Very Good” *non*-equilibrium which has even better payoffs. Subjects tended to follow the suggestions that implemented the good correlated equilibrium, but not those for the bad correlated equilibrium or the very good *non*-equilibrium.

Duffy, Lai and Lim (2017) focus on questions of learning to coordinate.

They compare direct recommendations to indirect recommendations—symbols with no pre-existing meaning—to no recommendations at all. They do this for two matching protocols, the first being random matching with a different partner on each play and the second being fixed matching with the same partner over repeated plays. The game is a Battle of the Sexes game, with the correlated equilibrium targeted placing probability one-half on each of the pure Nash equilibria.

With random matching, players used direct recommendations to correlate, as in previous studies. They improved their payoffs over those who had indirect recommendations or no recommendations. With fixed partners, they had an opportunity to learn to use indirect recommendations to coordinate. But instead, players learned to ignore all recommendations and coordinate by *taking turns*, as discussed in Vanderschraaf (2001) and Vanderschraaf and Skyrms (2003). In this way, they achieved payoffs equal to the expected payoffs of the correlated equilibrium, without experiencing the vicissitudes of chance.

All in all, there is robust evidence for the implementability of correlated equilibria. On the other hand, the experimental evidence on coarse correlated equilibria is thin, and what there is, is negative. Georgalos and SenGupta (2020) use a game with a unique Nash equilibrium in pure strategies. This equilibrium can be reached by iterated elimination of strictly dominated strategies. They test this Nash equilibrium against a coarse correlated equilibrium which delivers better expected payoffs. The coarse correlated equilibrium gives equal probability to the Nash and two strategy profiles symmetric about the Nash. Players can choose to commit to the coarse correlated equilibrium, in which case a computer carries out this strategy. Alternatively, they could elect to choose for themselves. A large majority elected to choose for themselves, and then played the Nash equilibrium. When only one of the two players chose to commit to the coarse correlated equilibrium, the other was shown the recommendations of the

coarse correlated equilibrium and then allowed to choose the action. In this case, those who chose for themselves tended to choose their end of the Nash equilibrium whatever the recommendation. Various variations and robustness checks did not change the main conclusion. Subjects did not choose the coarse correlated equilibrium.

One might wish for additional experiments, matching the variety that exist for correlated equilibria. The dominance-solvable Nash equilibrium may be too nice. The authors suggest in closing that coarse correlated equilibrium might do better if tested against Nash in a game where the unique Nash is in mixed strategies. The last example of Sect. 1 would be a possibility. If we are interested in spontaneous evolution of social conventions, we might like to see long series of trials without explicit recommendations, although this might make establishment of a coarse correlated equilibrium even harder. At present, such experiments have not yet been done.

6 Quasi-Conventions

When we extend the notion of convention from Strict Correlated equilibria to Strict Coarse Correlated equilibria, we arrive at the notion of a *Quasi-Convention*. With *Quasi-Conventions* we cast a wider net. In some situations, there are *Quasi-Conventions* that are better for all involved than any Correlated convention. (There are also ones that are worse.) The price we pay is that we generate strains of commitment.

Considering a society with widespread and persistent irrationality opens a Pandora's box, and calls the explanatory value of any of these equilibrium concepts into serious question. We will confine ourselves to considering a society of individuals who are at least approximately Bayesian rational. In such a society, there is no special attachment to commitment unless doing so maximizes expected utility. Quasi-conventions that are not correlated conventions are unstable. Correlated conventions are stable. Vanderschraaf gives the right explication of convention.

This does not mean that the notion of quasi-convention is useless. There is the possibility of embedding a given game in a larger game such that individuals following a quasi-convention, or approximating a quasi-convention, in the embedded game are maximizing expected utility in the context of the larger game.

Of course, embedding in a larger game can cover a multitude of sins, as the well-known folk theorems for repeated games show quite clearly. For *Quasi-Conventions* to be a useful category, the larger game must be of a particular kind, in which a norm of commitment to abide by the results of a lottery is sustained. Some real-world situations may be well-modeled in this way. In such a context, the concept of a *Quasi-Convention* can have a useful subsidiary role to play. *Quasi-Conventions* can be sub-conventions, modules that form part of larger genuine conventions.

Declarations

Conflict of interest I have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long

as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Aumann, R. (1974). Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1, 67–96.
- Aumann, R. (1987). Correlated equilibrium as an expression of bayesian rationality. *Econometrica*, 55, 1–18.
- Awaya, Y., & Krishna, V. (2019). Communication and cooperation in repeated games. *Theoretical Economics*, 14, 513–553.
- Barman, S., & Ligett, K. (2015). Finding any nontrivial coarse correlated equilibrium is hard. In *ACM conference on economics and computation*, EC '15. ACM.
- Binmore, K. (1994). *Game theory and the social contract, Volume 1: Playing fair*. MIT Press.
- Binmore, K. (1998). *Game theory and the social contract, Volume 2: Just playing*. MIT Press.
- Binmore, K. (2008). *Game theory: A very short introduction*. Oxford University Press.
- Blum, A., & Mansour, Y. (2007). From external to internal regret. *Journal of Machine Learning Research*, 8, 1307–1324.
- Bratman, M. (1987). *Intentions, plans, and practical reason*. Harvard University Press.
- Brown, G. W. (1951). Iterative solutions of games by fictitious play. In T. C. Koopmans (Ed.), *Activity analysis of production and allocation*. Wiley.
- Buchak, L. (2013). *Risk and rationality*. Oxford University Press.
- Cason, T., & Sharma, T. (2007). Recommended play and correlated equilibria: An experimental study. *Economic Theory*, 33, 11–27.
- Duffy, J., & Feltovich, N. (2010). Correlated equilibria, good and bad: An experimental study. *International Economic Review*, 51, 701–721.
- Duffy, J., Lai, E.K., & Lim, W. (2017). Coordination via correlation: An experimental study. *Economic Theory*, 64(2), 64, 265–304. <https://doi.org/10.1007/s00199-016-0998-8>
- Foster, D. P., & Vohra, R. (1997). Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21, 40–55.
- Fudenberg, D., & Levine, D. (1998). *Theory of learning in games*. MIT Press.
- Gauthier, D. (1986). *Morals by agreement*. Clarendon Press.
- Georgalos, K., Ray, I., & SenGupta, S. (2020). Nash versus coarse correlation. *Experimental Economics*, 23, 1178–1204.
- Hart, S. (2005). Adaptive heuristics. *Econometrica*, 73, 1401–1430.
- Hart, S., & Mansour, Y. (2010). How long to equilibrium? The communication complexity of uncoupled equilibrium procedures. *Games and Economic Behavior*, 69, 107–126.
- Hart, S., & Mas-Colell, A. (2003). Uncoupled dynamics do not lead to Nash equilibrium. *American Economic Review*, 93, 1830–1836.
- Hart, S., & Mas-Colell, A. (2006). Stochastic uncoupled dynamics and Nash equilibrium. *Games and Economic Behavior*, 57, 286–303.
- Hart, S., & Mas-Colell, A. (2000). A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68, 1127–1150.
- Huttegger, S. (2017). *The rational foundations of probabilistic learning*. Cambridge University Press.
- Lewis, D. K. (1969). *Convention: A philosophical study*. Harvard University Press.
- McClellan, E. F. (1990). *Rationality and dynamic choice: foundational explorations*. Cambridge University Press.
- Moreno, D., & Wooders, J. (1998). An experimental study of communication and coordination in noncooperative games. *Games and Economic Behavior*, 24, 47–76.

- Moulin, H., Ray, I., & Sen Gupta, S. (2014). Improving Nash by coarse correlation. *Journal of Economic Theory*, 150, 852–865.
- Moulin, H., & Vial, J. P. (1978). Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7, 201–221.
- Nisan, N., Roughgarden, T., Tardos, E., & Vazirani, V. (Eds.). (2007). *Algorithmic game theory*. Cambridge University Press.
- Rawls, J. (1971). *A theory of justice*. Harvard University Press.
- Robinson, J. (1951). An iterative method of solving a game. *Annals of Mathematics*, 54, 296–301.
- Roughgarden, T. (2016). *Twenty lectures on algorithmic game theory*. Cambridge University Press.
- Shamma, J., & Arslan, G. (2005). Dynamic fictitious play, dynamic gradient play, and distributed convergence to Nash equilibrium. *IEEE Transactions of Automatic Control*, 50, 312–327.
- Shapley, L. (1964). Some topics in two-person games. In M. Dresher, L. S. Shapley, & A. W. Tucker (Eds.), *Advances in game theory*. Princeton University Press.
- Strotz, R. H. (1955–1956) Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies*, 23, 165–180.
- Vanderschraaf, P. (1995). Convention as correlated equilibrium. *Erkenntnis*, 42, 65–87.
- Vanderschraaf, P. (1998). Knowledge, equilibrium and convention. *Erkenntnis*, 49, 337–369.
- Vanderschraaf, P. (2001). *Learning and coordination: Inductive deliberation, equilibrium, and convention*. Routledge.
- Vanderschraaf, P. (2019). *Strategic justice: Convention and problems of balancing divergent interests*. Oxford University Press.
- Vanderschraaf, P., & Skyrms, B. (2003). Learning to take turns. *Erkenntnis*, 59, 311–348.
- Young, H. P. (2004). *Strategic learning and its limits*. Oxford University Press.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.