

UC Irvine

UC Irvine Previously Published Works

Title

Invention and Evolution of Correlated Conventions

Permalink

<https://escholarship.org/uc/item/43n0g6xg>

Journal

The British Journal for the Philosophy of Science, 76(1)

ISSN

0007-0882

Authors

Herrmann, Daniel A
Skyrms, Brian

Publication Date

2025-03-01

DOI

10.1086/717161

Peer reviewed

Invention and Evolution of Correlated Conventions

Daniel A. Herrmann and Brian Skyrms

An important feature of many conventions is that the agents use an asymmetry to coordinate their behaviour. We call these ‘correlated conventions’. However, a puzzle arises: since any asymmetry works as well as any other, what are the relevant asymmetries on which a given population finds its correlated conventions? In order to gain traction on this question we need an account of both the invention and evolution of correlated conventions. Invention has remained largely unexplored in the literature. In this article we provide a simple model of the origin and subsequent dynamics of correlated conventions. This model can serve as a base for future investigation.

1. Introduction

Consider the animal that must decide whether to fight another for control over its territory. Or, consider the person who must decide whether to call a friend back after being disconnected from their call. In both strategic situations it is natural to think that the agent decides whether to take one action (be aggressive/call) or another (be passive/wait) based on an asymmetry (for example, ownership, or who called whom). This kind of strategy is a type of conditional strategy, in which the agent takes one action or another based on some kind of cue. When a conditional strategy governs interactions between organisms (as opposed to when it governs responses to nature) we call this strategy an ‘interactive strategy’. Informally, when such an interactive strategy has become fixed in the population, we call this a ‘correlated convention’.

One standard account of conventions is Lewis’s ([1969]), in which a convention is required to be a strict Nash equilibrium.¹ He gives the example of telephone calls in his hometown of Oberlin, Ohio. Because of some strange circumstance, all local calls in this town were disconnected three minutes into any conversation. In order to continue

¹ With some other conditions.

the conversation after the disconnection, one person must call back, and the other must answer but not call back herself. The convention that emerged was that the original caller would call back while the other person waited. Since both people want to resume the call, and both are worse off if either deviates from the plan unilaterally, this is a convention in Lewis's sense.

An important feature of this particular convention is that the agents use an asymmetry to anti-coordinate their behaviour. This feature, present in many conventions, motivated Vanderschraaf ([1995]) to extend Lewis's account by replacing the Nash equilibrium requirement with that of a correlated equilibrium, as defined by Aumann ([1974]). Indeed, we can recast Lewis's Oberlin example as an instance of this kind of correlated convention, where players condition their behaviour on the caller versus called asymmetry.

However, a puzzle arises. Any asymmetry works as well as any other. Indeed, Lewis ([1969], p. 43) writes: 'Other regularities [besides the caller calling back regularity] might have done almost as well. It could have been the called party who always called back, or the alphabetically first, or even the older. Any of these regularities could have become the convention if enough of us had started conforming to it'. The observation is that there are many asymmetries on the basis of which a convention could be established.

Lewis's picture lacked a dynamic account of how we come to any particular convention; he was merely concerned with which were the possible conventions. However, even in a more dynamic evolutionary context, the puzzle can persist. Eshel ([2005]) notes that there are many asymmetries on which agents might condition their strategy. He writes: 'As a simple example, it is easy to see that in any Hawk–Dove conflict within a human society, a behavioral rule of giving priority to the contender with, say, longest thumb can well replace Ownership–Priority, demonstrated by Maynard Smith and Parker as a strict, asymmetric ESS [evolutionary stable strategy]. Moreover, Thumb–Priority would be as efficient as Ownership–Priority in preventing aggressive confrontations within the population' (Eshel [2005], p. 12). This leads to the question, why do we observe ownership priority as opposed to thumb priority? In general, what are the relevant asymmetries on which a given population founds its correlated conventions? The same puzzle of asymmetry selection occurs in both Lewis's philosophical account of conventions and in evolution game theory.

Eshel emphasizes how serious a challenge this plurality of asymmetries poses for evolutionary game theory's ability to shed insight on the behavioural rules in populations. In his words: 'given the high dimensionality of asymmetry in natural populations, *none of the behavioral rules, observed in such populations, can be explained on the pure basis of population game theory*, without resorting to further information about historical facts and evolutionary paths' (Eshel [2005], p. 13). When trying to model these kinds of correlated conventions the standard approach in population game theory is that the modellers identify a particular asymmetry—say, ownership—and incorporate that asymmetry into the model. We briefly review

some of this literature in section 2, as it is an essential piece of the puzzle, and sheds light on particular dynamics we observe in nature.

The question Eshel identifies, and the one at which our article aims, is different. Instead of investigating a question like ‘why the ownership-priority convention as opposed to the opposite convention?’, we investigate the question ‘how does the set of strategies over which selection is operating arise in the first place?’.

If we want a more general account of the development of these kinds of correlated conventions, both of these questions will need to be addressed. To address the second question, we need to have a better understanding of the path-dependency and historical dependence that Eshel identifies. To this end it would be very useful to have a precise model of the invention and dynamics of correlated conventions. In particular, it should be one in which correlated conventions themselves can evolve without assuming which asymmetries are and are not in place, but rather by letting arbitrary asymmetries be introduced via invention.

Invention has remained largely unexplored in most relevant models of evolutionary game theory²; a richer approach is needed. In this article we take the first steps towards developing such an approach by giving a simple model of the origin and subsequent dynamics of correlated conventions. This model can serve as a base for future investigation into the puzzles surrounding correlated conventions. In particular, analyses that seek to answer the first question might be layered on top of the base model we provide. We leave such analyses for future work.

Our model incorporates aspects of and draws inspiration from various others in the literature. After briefly introducing correlated conventions and associated puzzles, we review these models. We then propose our new model, a form of reinforcement with invention,³ and discuss both analytic and simulation results. The results help us address three main questions: Do correlated conventions always take over? When they do take over, how fast? Can alternative conventions co-exist? We also discuss the connection between order of invention and salience, and how this plays a role in establishing convention. We end with some concluding remarks.

2. Correlated Conventions

Vanderschraaf’s ([1995]) extension of Lewis’s account of conventions uses Aumann’s ([1974]) idea of a correlated equilibrium. In the contexts Aumann considers, players of a game observe the value of an external random variable on which they can condition their acts. This extends the set of possible strategies the agent might adopt, and thus opens up new possibilities for coordination. A joint correlated strategy tells each player what act to take based on what outcome they observe. A

² A notable exception is a recent paper by Morsky and Akçay ([2019]), the model of which is quite different from the one we provide here.

³ Alexander et al. ([2012]) previously used a similar approach to model the invention of new signals in signalling games.

joint correlated strategy is a correlated equilibrium if, no matter what each player observes, all players want to follow the recommendation of the joint correlated strategy. Importantly, correlated equilibria can allow for a higher social payoff. For a clear discussion of this and other properties of correlated equilibria, see (Aumann [1974], [1987]).

As an example of a correlated equilibrium, consider two players in the hawk–dove game. This is an anti-coordination game: whatever strategy you choose, I want to do the opposite, and vice versa. If we do not allow conditional strategies, there are two pure Nash equilibria in this game, and one mixed Nash equilibrium. If we introduce a random variable on which the two players can condition their strategy, this introduces correlated equilibria. Suppose this variable can take on two values, zero or one. Consider the following joint strategy in which one player is aggressive if the variable is one and passive if it is zero, and the other player is aggressive if the variable is zero and passive if it is one.

Neither player has any incentive to deviate from this joint strategy. If you are the first player and the random variable takes on the value zero, even though you would prefer it if the variable had been a one, you will still follow the strategy and be passive since you expect me to be aggressive. If it is a one, then you prefer to be aggressive when I am passive. This is a correlated equilibrium. The joint strategy ensures that the players successfully anti-coordinate.

Maynard Smith and Parker ([1976]) consider the same example in the context of evolutionary game theory. They take the hawk–dove game to be a model of resource competition. They consider a population game with only one population in which there are hawks, which are always aggressive, and doves, which are always passive. In this case the evolutionary dynamics drives the population to a mixed state,⁴ mirroring the mixed Nash equilibrium in the static game: when there are mostly hawks in the population it is better to be a dove, and vice versa. The proportion of hawks and doves at which each type has an equal expected payoff is a mixture. The details depend on the actual payoffs of the game.

Maynard Smith and Parker also consider the possibility that individuals might condition their behaviour on an external random variable. They consider an instance of this in the hawk–dove game, in which one individual is always an intruder of a territory and another is the resident. The intuitive strategy ‘hawk if resident, dove if intruder’ is an equilibrium. The asymmetry of ownership allows the individuals to successfully anti-coordinate. Following Vanderschraaf, we can say that a population in which everyone follows this strategy has evolved a correlated convention. This is the ownership, or bourgeois, convention.

This convention is also an instance of what Maynard Smith and Price ([1973]) called an ESS. Informally, an ESS is a strategy that, once established in the population,

⁴ Other examples of populations driven to mixed states include sex ratios (Hamilton [1967]) and fair resource division in simplified Nash.

cannot be invaded by another strategy.⁵ The ESS concept was introduced as an attempt to identify candidates for stable outcomes of strategic interactions within animal populations. Initially this definition required that individuals have an equal probability of encountering any other individual in the population. However, others such as Bomze ([1986]) and Vickers and Canning ([1987]) showed how this assumption could be relaxed.

Allowing asymmetries between agents in the population has dramatic effects on the possible ESS in a game.⁶ A rather striking result by Reinhard ([1980]) illustrates this; motivating his result with the bourgeois convention, he shows that any ESS of an asymmetric population game must be a pure strategy. Even small asymmetries between players can have major effects on the outcome.

Eshel ([2005]) reminds us that Maynard Smith and Parker ([1976]) had noted in their paper that the opposite of the bourgeois convention ('hawk if intruder, dove if owner') works just as well to create a correlated equilibrium (they call this the paradoxical strategy).⁷ The upshot is that when there is an asymmetry, either way players condition their strategy on the asymmetry works.

There is a rich literature exploring how specific properties of particular asymmetries lead to conventions. In particular, there are many sophisticated proposals for why something like the bourgeois convention is more likely to evolve than the paradoxical convention. Maynard Smith and Parker ([1976]) note that the 'hawk if owner, dove if intruder' strategy is prevalent in nature, and suggest that this arises because in many cases the pay-off for victory in a conflict is slightly higher for the owner of a resource (see also Maynard Smith [1982]).⁸ Maynard Smith and Parker ([1976]) also discuss the case in which larger individuals have a higher probability of winning an altercation, and derive conditions under which agents will respect body-size priority (Eshel [2005] also discusses this case). Maynard Smith ([1982]) develops the discussion of these asymmetric payoff cases further. He notes that if hawks and doves are the only strategies in the population, then there is a mixed equilibrium that can't be invaded by hawks, doves, or the paradoxical strategy, but can be invaded by the bourgeois strategy (Maynard Smith [1982], p. 105). Barton ([1979]) provides an analysis in which geographical constraints play a key role in determining which conventions can coexist. Barton and Hewitt ([1981]) develop this idea further and apply it to questions of speciation.

⁵ The formal definition of an ESS is as follows. Let $\pi_{\sigma,\tau}$ be the average payoff that (possibly mixed) strategy σ gets when playing against strategy τ . Then we say that σ is an ESS if and only if, for any mutant τ , $\pi_{\sigma,\sigma} \geq \pi_{\tau,\sigma}$ and, if $\pi_{\sigma,\sigma} = \pi_{\tau,\sigma}$ then $\pi_{\sigma,\tau} > \pi_{\tau,\tau}$.

⁶ This is an example of the kind of situation discussed above, in which the set of possible strategies agents might adopt is extended by allowing agents to condition their actions on random variables. In this context the random variable is an asymmetry.

⁷ For a review of the evolution of respect for property, see (Sherratt and Mesterton-Gibbons [2015]).

⁸ For example, Maynard Smith ([1982], p. 101) writes, 'Thus the value of a territory may be greater to an owner, who has already learnt about the distribution of food, refuges, etc. In some cases ownership may confer advantages in an escalated contest'. This type of explanation emphasizes that differences in pay-offs of the initial interactions in a population can drive a population to one convention or another. Sugden ([2004]) draws similar lessons for experiential learning in economics.

Another important factor in determining which conventions are likely to arise is salience. The core idea is that in many cases there are going to be certain properties of an asymmetry (or, more generally, some property of a particular Nash equilibrium) that make that asymmetry likely to be the property on which a convention is founded. Inspired by Schelling's ([2004]) observation that agents are often able to coordinate their acts by exploiting certain salient features of a decision problem, Lewis ([1969]) appealed to salience in order to explain how certain conventions arise. There has been much further development in this area. Mehta et al. ([1994a], [1994b]) have done thorough empirical studies showing that humans are able to coordinate their acts using salient features of the decision problem. Sugden ([1995], [2004], [2011]), Goyal and Janssen ([1996]), Schlicht ([1998]), Mesterton-Gibbons and Adams ([2003]), Alberti et al. ([2012]), and LaCroix ([2020]) have also provided explanations of the emergence of convention that feature salience.⁹

Once our attention is drawn to the dramatic effects of asymmetries, we are faced with important new questions, many of which are laid out forcefully by Eshel ([2005]). In particular, as discussed in section 1, Eshel notes that there are many asymmetries on which agents might condition their strategy. With an account of evolutionary dynamics in the foreground, Eshel takes the upshot of this observation of many different asymmetries to be that we should expect the dynamics of what we here call correlated conventions to be very path-dependent. In particular, the problem that we would need to address is, what is the set of asymmetries over which selection is operating? This severely limits the usefulness of the ESS concept.

There is a certain irony that Eshel identifies. Maynard Smith and Price's introduction of the idea of the ESS was intended to be making progress towards identifying the possible states in which actual populations might end up. However, with Maynard Smith and Parker's observation that the opposite of the bourgeois strategy works just as well, they (followed later by Eshel) also took the initial steps towards showing how limited an approach this is.¹⁰

3. Previous Work

3.1. Reinforcement and replication

In reinforcement learning an agent is faced with a sequence of identical decision problems. At each stage she must make a decision from a finite number of possible actions, a_1, \dots, a_k . Each time the agent makes a decision she chooses an action with a probability proportional to the rewards she has accumulated from the past instances she chose that action.¹¹

⁹ The order of invention in the model we provide can be interpreted as reflecting a salience ordering of different asymmetries. We discuss this in section 4.

¹⁰ Mohseni ([2019]) discusses other limitations of equilibrium concepts in evolutionary game theory.

¹¹ Pemantle ([2007]) provides a survey of reinforcement learning processes.

A simple model that captures the logic of reinforcement learning is an urn model. To see how we can represent this with an urn process, assign a different colour to each of the k actions the agent might take. The agent starts with an urn in which there is at least one ball of each colour. Each time the agent must choose an action she draws a ball from the urn at random and takes the action corresponding to the colour she observed. She experiences her reward, returns the ball she drew to the urn, and then adds an additional group of balls of that colour to the urn. The size of this group is proportional to the magnitude of the reward she experienced.

In the context of two-person games, Roth and Erev ([1995]; Erev and Roth [1998]) showed that this kind of learning dynamic can allow the players to successfully learn to play a host of games. Beggs ([2005]) and Hopkins and Posch ([2005]) have examined the long term convergence for this kind of reinforcement learning in specific classes of games. Argiento et al. ([2009]) have demonstrated convergence in a Lewis signalling game.¹² Barrett et al. ([2017]) modify the simple reinforcement learning in a Lewis signalling game by combining it with a form of trial-and-error learning.

Even though this urn model was first used to describe individual learning, we can observe the strong resemblance between reinforcement learning and evolution in a small finite population. Each individual in a population is like a ball, and the strategy each follows is like the action corresponding to the colour of the ball. Reproduction is the reinforcement mechanism.

Schreiber ([2001]) extends this idea with the addition of a game to the dynamics. An urn holds a set of balls representing the initial population, with a different colour corresponding to each one of the k types in the population. When a game with n players determines the reproductive fitness of individuals, reinforcement proceeds as follows: On each round, n balls are drawn at random from the urn, and the individuals they represent interact according to the game. The balls are returned to the urn, and additional balls of those colours are added to the urn proportional to the pay-offs of that round of play.

3.2. Invention

The Schreiber urn model gives us a way to think about reinforcing successful strategies in the population. However, in order for the strategy to be reinforced it must already be present in the population. How about mutation, or invention?

Models of neutral evolution, in which types have no effect on fitness, can provide a jumping off point for invention. Hoppe ([1984], [1987]) provides an urn model that

¹² The Lewis signalling game is a game that models how two players might use arbitrary signals to transmit information. Lewis ([1969]) introduced the signalling game to help explain how communication could arise without there being a pre-existing language in place. Skyrms ([2010]) developed this analysis further by investigating signalling games in evolutionary and learning contexts.

allows for the genesis of new types.¹³ The model works as follows: The urn starts out with various balls of different colours, including at least one black ball: this is the mutator ball. A ball is drawn at random from the urn in each round. If the ball is not black, it is replaced in the urn along with an additional ball of the same colour. If the black ball is drawn it is replaced in the urn as well, but instead of adding another black ball, we add a ball with a new colour not yet present in the urn. There are no limits on the number of new colours that might be added; the limiting result of this process has an infinite number of colours.

This is of course a model of neutral evolution; there is no fitness to track. Types are distinguishable from one another only based on their labels. We, however, are interested in invention in a context in which types have different fitness. Combining the Hoppe urn model with the Schreiber urn model by adding a mutator ball to the Schreiber urn would introduce a mechanism of invention to a model the dynamics of which are indeed sensitive to fitness.

An example of adding this kind of invention to a model is the work of Alexander et al. ([2012]), who consider a quite natural way to incorporate invention into Lewis signalling games. They note that if the number of signals, m , is less than the number of states of nature, n , then it seems reasonable that the sender might invent new signals.¹⁴ Furthermore, it is clear in this context that the actual signal used doesn't have any payoff consequences; it only matters that the sender and the receiver coordinate their behaviour.

We pursue a similar strategy in section 4 by having the payoffs of new strategies depend on a base game that doesn't itself change. In our limited context this allows us to introduce new strategies in a coherent way.

4. Inventing Correlated Conventions: The Model

Now that we have all the moving pieces before us we can describe a simple model of the invention and evolution of correlated conventions. We take the approach sketched in section 3.2 and use the Schreiber urn model of evolution with the addition of mutator balls. When a mutator ball is drawn, a new type of strategy, one in which acts are conditions on an external random variable, is added.

Schreiber ([2001]) also allows for death in his model. When an individual dies, their corresponding ball is removed from the urn. The probability that an individual dies can be chosen to depend on the outcome of the games they play. Death introduces the possibility of extinction, both for types and for the whole population.¹⁵

¹³ Although Hoppe introduced his urn model as a model of neutral evolution, the model also has rich connections to Bayesian inference and inductive logic, as well as artificial intelligence. For a discussion of generalizations of this model and its connections to these fields, see (Fortini and Petrone [2012]).

¹⁴ To introduce invention, instead of the sender's urns starting with a selection of signals she might send, they instead start with just a mutator ball in each. The dynamics are the same as normal, except when the mutator is drawn a new colour (corresponding to a new signal) is added to the urn.

¹⁵ Death for population resembles forgetting for individuals. Roth and Erev notice the possibility of forgetting in their urn model; Barrett and Zollman ([2009]) further discuss forgetting in this context.

The model we consider here will not include any death mechanics. However, this could be investigated in future work.

Adding new types to the model raises a whole host of questions. How are these new types generated? How do they interact with the old types? What are the payoff consequences on the invented strategy when it interacts with old types?

Drawing inspiration from the approach of Alexander et al. ([2012]), in order to solve these problems our model uses a base game (for example, hawk–dove). All of the payoff consequences of new types will be mediated through the actions available in this base game. The conditional strategies introduced by mutation or invention are of the form ‘if the random variable takes on value x , take action a ’ (specified for each value the random variable might take on) where a is one of the actions available to agents in the base game.¹⁶ Just as new signals are unproblematic in the context of invention in signalling games since all that matters is whether the agents can successfully coordinate their behaviour, new conditional strategies are unproblematic in the context of our model because they only lead to actions already available in the base game, for which all payoff details are fully specified. Other complexities of introducing a random variable are avoided because the random variables we consider are payoff-neutral. This would not be the case if we considered games in which an asymmetry seemed relevant (perhaps relative strength or size in a game that involved a physical altercation, for example).

However, there will be some choices that have to be made: What is the nature of the random variables? Are they cyclical, like day or night? Are they multi-valued, like temperature? Are they asymmetric, like the random variable Eshel’s thumb-priority rule uses? Different choices would lead to differences of the model.

The final thing we need to consider is the order of invention. We could think of the new strategies being introduced in the population as being generated where the particular random variable on which it conditions is randomly chosen; this allows us to incorporate the idea of an historical accident of attention (an example of Eshel’s historical fact) into the model. Or, instead, inspired by the literature we reviewed in section 2, we could think of the order of invention as reflecting some kind of natural salience ordering, where some random variables are more salient than others. This salience ordering would be relative to the intended context of application. This would then let us ask and answer questions such as, how path-dependent is the process? If a strategy is invented earlier, how much more likely is it to become a correlated convention? And so on. We are proposing a general framework in which a particular base game, type of random variable, and order of invention would need to be further specified. In the remainder of the article we examine one simple instance of this general framework.

¹⁶ In the particular model we describe in section 5, the single random variable on which agents will condition is a function of one property from each of the individuals in the interaction.

5. Model Example

Our example is a model of evolution in a finite, growing population. We take the particular instance of the hawk–dove game in table 1 as our base game. There is a countably infinite list of visible traits. For each agent i and each trait j , nature assigns a random value $x_{i,j} \in (0, 1)$ representing the j th trait value of the i th agent.¹⁷ Example traits are thumb length, plumage brightness, and so on. Any time an agent is added to the population, either through mutation or through reinforcement, nature assigns a random value to each trait for that agent.¹⁸

The population starts in a state in which the only strategies agents play are the unconditional ‘always hawk’ or ‘always dove’. The dynamics proceed similarly to the Schreiber urn model. At each time step a fixed proportion of agents are drawn from the population and paired off at random. Each pair plays the hawk–dove game, and replicate according to their payoffs.

The main difference from the Schreiber urn model occurs when a mutator ball is drawn.¹⁹ In this case a mutant is added to the population. There are two mutator balls, one of which we will call the big mutator ball and the other the small mutator ball. The i th time the big mutator ball is drawn a mutant with the conditional strategy ‘hawk if my i th property is larger than my opponent’s i th property, dove otherwise’ is added to the population. Similarly, the i th time the small mutator ball is drawn a mutant with the strategy ‘hawk if my i th property is smaller than my opponent’s i th property, dove otherwise’ is added to the population.²⁰ Having both mutator balls allows us to use the model to explore how the contradictory strategies, Big i (B_i) and Small i (S_i), interact.

6. Equilibria Analysis

The model has a potentially infinite number of possible strategies, since for each $n \in N$ both B_n and S_n are strategies that might be invented at some point. However at each point in time only a finite number of strategies will have been invented. We can construct a game matrix to capture the payoffs to different types in the population based on the base game. We illustrate this in table 2, with an example where only hawks, doves, B_1 , S_1 , B_2 , and S_2 are in the population. The top left corner of the payoff matrix reflects the base hawk–dove game. When a hawk encounters a conditional strategy, since on average half the time the conditional strategy will be aggressive

¹⁷ This aspect of our model is similar to how Axelrod ([1997]) handles features of different cultures in a model of the dissemination of culture, although he considers only a finite set of features and values they can take.

¹⁸ Specifically, the traits are drawn from a uniform distribution over the unit interval. Note that since each draw is independent, the different traits will be uncorrelated with one another, and with the strategy the agent uses. Future work could relax any of these assumptions.

¹⁹ Specifically, as with the Schreiber urn model, the new agents inherit the strategies. Note that in our model, agents do not inherit the traits: as stated above, those are assigned randomly whenever a new individual is added to the population.

²⁰ One might wonder what happens if the properties are exactly equal. The first thing to note is that this is a probably zero event, since the properties are drawn according to a uniform distribution on $(0, 1)$. For concreteness in the simulations, which only approximate a probability density function, we used ‘ \geq ’ for the strategies generated from the big mutator ball, and ‘ $<$ ’ for the strategies generated from the small mutator ball.

Table 1. An example of a hawk–dove normal-form game.

	Hawk	Dove
Hawk	0, 0	3, 1
Dove	1, 3	2, 2

and half the time it will be passive,²¹ the hawk gets an average payoff of 1.5, and the conditional player gets an average payoff of 0.5. Similarly, when a dove encounters a conditional strategy player, the dove gets an average payoff of 1.5 and the conditional strategy gets an average payoff of 2.5.

When a conditional strategy player encounters a player using the same conditional strategy, they perfectly anti-coordinate their strategies. Thus the average payoff to a player in this kind of encounter is 2. When a conditional strategy player encounters a player using the contradictory strategy—for example, B_2 plays against S_2 —they have a 0.5 probability to both play aggressive and a 0.5 probability to both play passive. Thus their average payoff is one. When two players that condition on different variables encounter each other there is a 1/4 probability that their play ends in any of the action pairs; thus their average payoff against each other is 1.5.

Equilibria in this game must be symmetric since players do not condition on whether they are the row player or the column player. There are four categories of Nash equilibria in this game.

The first is the mixed Nash equilibrium from the base hawk–dove game.²² With these specific payoffs the equilibrium is one in which players play dove 1/2 of the time and hawk 1/2 of the time. This equilibrium has an expected payoff of 1.5.

The second category we call the ‘mixed conventions’. A mixed convention equilibrium is a mixture of conditional strategies that condition on different properties. In this game there are four mixed convention equilibria:

$$0.5 \cdot B_1 + 0.5 \cdot B_2,$$

$$0.5 \cdot B_1 + 0.5 \cdot S_2,$$

$$0.5 \cdot S_1 + 0.5 \cdot B_2,$$

$$0.5 \cdot S_1 + 0.5 \cdot S_2.$$

Each of these equilibria has an expected payoff of 1.75.

²¹ Recall that the player playing the conditional strategy is not using a mixed strategy. Rather, since half the time the player’s property on which it conditions its behaviour, say x_i , will have a value greater than its opponent’s x_j , it will play aggressive half the time and passive half the time.

²² Recall this is a single population game. Hence probabilities in the mixed strategy Nash equilibrium can be interpreted as population proportions of types.

Table 2. Extended game with two properties.

	Hawk	Dove	B_1	S_1	B_2	S_2
Hawk	0, 0	3, 1	1.5, 0.5	1.5, 0.5	1.5, 0.5	1.5, 0.5
Dove	1, 3	2, 2	1.5, 2.5	1.5, 2.5	1.5, 2.5	1.5, 2.5
B_1	0.5, 1.5	2.5, 1.5	2, 2	1, 1	1.5, 1.5	1.5, 1.5
S_1	0.5, 1.5	2.5, 1.5	1, 1	2, 2	1.5, 1.5	1.5, 1.5
B_2	0.5, 1.5	2.5, 1.5	1.5, 1.5	1.5, 1.5	2, 2	1, 1
S_2	0.5, 1.5	2.5, 1.5	1.5, 1.5	1.5, 1.5	1, 1	2, 2

The third category are the contradictory equilibria. They are of the form

$$a \cdot B_1 + a \cdot S_1 + b \cdot B_2 + b \cdot S_2 + c \cdot \text{hawk} + c \cdot \text{dove},$$

where $a, b, c \geq 0$, either $a > 0$ or $b > 0$, and $2a + 2b + 2c = 1$. All of these have an expected payoff of 1.5.

Finally, there are the pure convention equilibria. These are equilibria in which everyone plays the same conditional strategy. There are four of these in this game, one for each conditional strategy. The average payoff for a player in this kind of equilibrium is 2.

The type of equilibrium with the highest payoff is the pure convention equilibrium with a payoff of 2. Thus we see that introducing conditional strategies can improve the average payoff for individuals in the population.²³ However, in the dynamic context we don't just want to know what the Nash equilibria are—we also want to know which of these equilibria are stable. Of the four types of equilibria described above only the pure convention is an ESS. It is easy to see that each pure convention equilibrium is an ESS, since each is also a strict Nash equilibrium.

The kind of analysis above holds when there are more than two properties for which conventions have been invented. For each possible set of strategies at play in the game, the only ESS are pure convention states. Furthermore, since each strategy that hasn't yet been invented interacts with each of the already present strategies in the same way, there are no problems that come with the introduction of new strategies for the qualitative feature of the analysis.

This analysis vindicates Eshel's concern—among so many different possible conventions a population might follow, why the particular one on which it settles? Eshel ([2005], p. 12) notes that this problem is somewhat hidden when one 'arbitrarily concentrates on one single parameter of asymmetry'. When one considers multiple asymmetries (for example, Eshel and Sansone [2001]) the situation is less straightforward. Our model, which allows for a possibly infinite number of different asymmetries, makes very clear that the standard analysis focusing on an artificially chosen few is insufficient; each of an infinity of possible asymmetries can be used to form stable correlated conventions, and each gives the same average payoff for members

²³ As emphasized by Maynard Smith.

following that correlated convention. In order to shed more light on this problem we need to look at the dynamic picture.

7. Simulation Results

We ran simulations using the dynamics described in section 5. Our model can accommodate an infinite number of possible asymmetries. The initial state of the population was one hawk and one dove, as well as the two mutator balls. An individual is drawn at random from the population each round. If the individual is a mutator, then a new type is added as described in section 5. If the individual is not a mutator ball, then 1% of the population is drawn at random,²⁴ paired up to play the game, and types reinforced according to the outcome of the interactions.

For $n \in \{1, \dots, 6\}$ we ran 500 simulations out to $n \cdot 100$ rounds. If a single conditional strategy made up more than 75% of the population at the end of the simulation then we called that a ‘convention takeover’.²⁵ Similarly, if hawks and doves together made up more than 75% of the population then we called that a ‘hawk–dove takeover’. The results are shown in figure 1.

At the end of 100 rounds, only 5.6% of populations end up at a state in which a single conditional strategy makes up more than 75% of the population. Furthermore, more than half of populations end up at a state in which hawks and doves together have taken over.²⁶ However, as the number of rounds played increases, the proportion of runs that end in a convention takeover increases. By 600 rounds, 26.2% of runs ended with a convention takeover. Although we were unable to run simulations past 600 rounds because of computational limitations, it seems that the longer we allow the population to evolve the more likely it is for a convention to take over. We thus have our first moral.

Moral 1: Invention of conditional strategies allows for new strategies to take over the population, establishing a convention. Furthermore, as the number of rounds increases so does the probability that a convention takes over.

There is another layer to this story. We aren’t only interested in how likely it is for a convention to take over—we are also interested in how well the population does once a convention has taken over. If only hawks and doves are in the population, then at the mixed Nash equilibrium from the base hawk–dove game the expected number of individuals added to the population from one interaction is three. However, at a pure convention equilibrium the expected number of individuals added to the population from one interaction is four. Thus we would expect populations

²⁴ If 1% of the population contains fewer than two individuals, then two individuals are drawn instead.

²⁵ The 75% threshold is somewhat arbitrary. However, very long simulations suggest that once a single conditional strategy makes up more than 75% of the population, it is very likely to continue growing in proportion.

²⁶ In these cases the relative proportions of hawks and doves is very close to one-to-one, as we would expect from the equilibrium analysis.




Figure 1. Fraction of runs that resulted in a convention or hawk–dove (HD) takeover. Note that the proportions do not sum to one, as there are cases where neither a convention nor the hawk–dove pair make up more than 75% of the population.

in which a convention takes over to be larger than those in which hawks and doves take over.

The simulation results bear this out. As seen in figure 2, the jump in population size of populations in which a convention took over from 500 to 600 rounds is quite severe, leaving the other populations far behind, and suggests that a convention taking over can dramatically increase the size of the population. Furthermore, even after only relatively few rounds the conventional populations enjoy a larger population size. Thus, it seems plausible that conventions could have a large effect in situations in which group selection is at play.²⁷ In particular, if there were competition between groups in which a larger group had an advantage, then we would expect the group with a convention to be more successful. Thus we have our second moral.

Moral 2: Populations in which a convention takes over have a large size advantage over those in which one does not. This may have implications for group selection.

We are also interested in the connection between the order of invention, which we can interpret as reflecting a salience ordering, and the establishment of a convention. One important question is: when a convention is successful, is it always based on the most salient asymmetry? Our simulation results answer this question in the negative. Indeed, although the first conditional strategy invented does enjoy a higher likelihood of becoming the dominant convention, 40.6% of conventions that take over were not the strategy invented first. Furthermore, 21.2% of conventions that successfully took over were invented third or later.²⁸ Thus, even though a convention that is

²⁷ For a discussion on group selection and its connection to individual selection, see (Wilson [1975]). For group selection in the context of the human behavioural sciences, see (Wilson and Sober [1994]).

²⁸ These results are from 300 simulations each lasting 400 rounds.

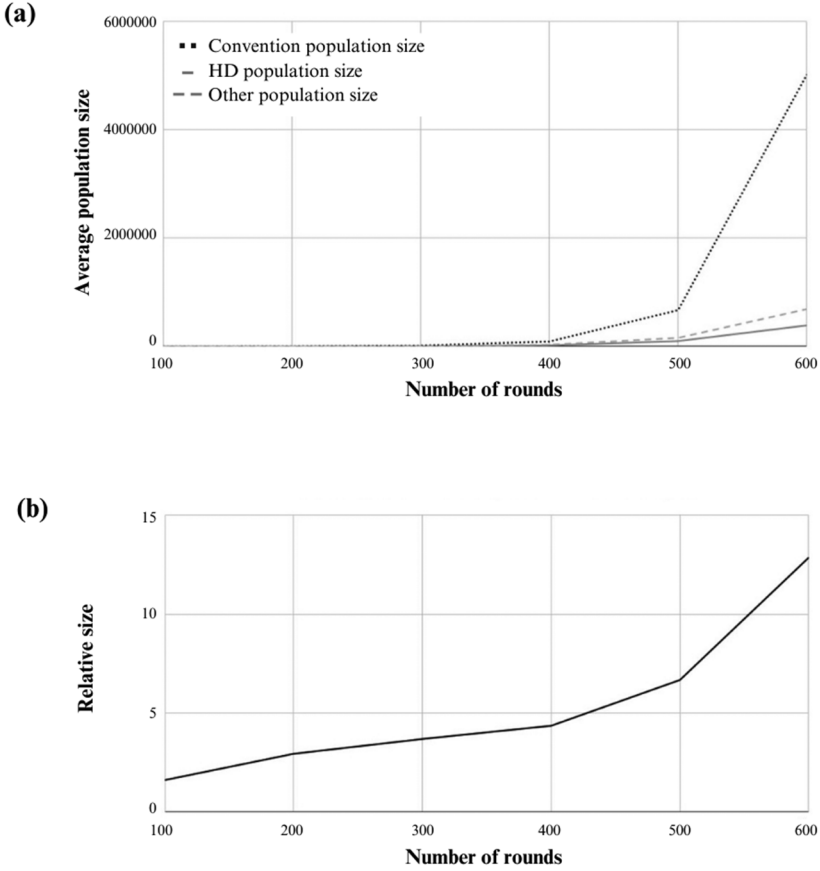


Figure 2. (a) The average population size of three different types of populations: those in which a convention took over, those in which hawks and doves took over, and those in which neither took over. (b) The average size of a population in which a convention took over divided by the average size of a population in which hawks and doves took over.

invented earlier is more likely to become the dominant convention, a non-trivial proportion of the successful conventions were not based on the most salient properties. Under the interpretation of order of invention reflected salience we have our third moral.

Moral 3: When correlated conventions do evolve they are more likely to be based around the most salient property, but this is far from guaranteed.

8. Conclusion

In this article we provided a model of a much neglected aspect of conventions: a dynamic story of their formation. This involved both aspects of invention and evolution.

The dynamic picture gives us the tools to tackle problems that arise from the multiplicity of exploitable asymmetries. By incorporating invention, we remove the arbitrary focus on only one asymmetry; our model respects the fact that any asymmetry can serve as the basis of a convention by letting arbitrary asymmetries be introduced via invention. This provides a direction for a possible solution to Eshel's challenge for population game theory.

We have opened up new avenues of inquiry by incorporating invention into evolutionary dynamics. We find that correlated conventions arise spontaneously in some cases, but that it is also common for this to fail. We find that when they do arise the population tends to be much larger than cases in which they do not, which may be relevant for contexts in which group selection is a relevant factor.

There are many possible ways to extend this work. These fall into two main categories: modifying the model and extending the analysis. Of course, it would be interesting to see both carried out in the same project. It is important to understand these possible extensions, because they also show the limitations of the current model.

Modifying the model: As with any model, there are many details of the model that need to be specified. We discussed a few of these in section 4; it would be informative to see how the qualitative results change (or not) given different choices of these parameters. In addition to those we mentioned there, following some of the work we discussed in section 2, it would be straightforward to incorporate asymmetries that lead to different payoffs. Consider body size as the asymmetry. Perhaps larger individuals have a higher probability of winning an altercation. Or perhaps when one animal has spent time in a territory (they are an owner) the territory is worth more to it than to one that hasn't spent much time there, since the former knows where the good shelter is. Both of these kinds of asymmetries in payoffs could be incorporated into the model and the effect they have on the outcomes explored.

The evolutionary dynamics are another natural feature to investigate. Our model built on the Schreiber ([2001]) urn model, which has important connections to reinforcement learning and, with the addition of the mutator balls, lends itself well to a context of invention. There are, of course, many other important dynamics present in the literature (for example, the replicator dynamics). It would be worth investigating how one might adapt the model to incorporate these dynamics, and how this would affect the results. Another significant change to the dynamics would be the addition of some kind of death process. We are interested in understanding whether and how this would change the qualitative results. There are reasons to expect that changes to the dynamics that affect population size might have interesting consequences for the probability that a convention takes over (Otto and Whitlock [1997]).

Extending the analysis: It would also be worthwhile to get a more fine-grained understanding of the dynamics of the recently invented conventions. For example, it would be informative to see how noisy the process of taking over is in the short run, and how many false starts there are. Additionally, it would be interesting to get a better sense of the diversity of strategies over time. In this vein, getting a better

sense of the rate at which a convention can take over, and connecting this up to evolutionary time-scales, would help us understand the extent to which this model can inform us about the world.

It would also be interesting to understand how the probability that a convention takes over changes as a population grows (extending fig. 1), and what the limiting properties of the process are. Some of this investigation could be carried out with the help of simulations, and some with analytic techniques.

Acknowledgements

We would like to thank Simon Huttegger, Jeffrey Barrett, Cailin O'Connor, Aydin Mohseni, and the three anonymous referees for helpful comments on earlier drafts. We would also like to thank the attendees of the social dynamics seminar at the University of California, Irvine, for feedback.

Daniel A. Herrmann
Department of Theoretical Philosophy
University of Groningen
Groningen, The Netherlands
d.a.herrmann@rug.nl

Brian Skyrms
Department of Logic and Philosophy of Science
University of California, Irvine
Irvine, CA, USA
bskyrms@uci.edu

References

- Alberti, F., Sugden, R. and Tsutsui, K. [2012]: 'Salience as an Emergent Property', *Journal of Economic Behavior and Organization*, **82**, pp. 379–94.
- Alexander, J. A., Skyrms, B. and Zabell, S. L. [2012]: 'Inventing New Signals', *Dynamic Games and Applications*, **12**, pp. 129–45.
- Argiento, R., Pemantle, R., Skyrms, B. and Volkov, S. [2009]: 'Learning to Signal: Analysis of a Micro-Level Reinforcement Model', *Stochastic Processes and Their Applications*, **119**, pp. 373–90.
- Aumann, R. J. [1974]: 'Subjectivity and Correlation in Randomized Strategies', *Journal of Mathematical Economics*, **1**, pp. 67–96.
- Aumann, R. J. [1987]: 'Correlated Equilibrium as an Expression of Bayesian Rationality', *Econometrica*, **55**, pp. 1–18.
- Axelrod, R. [1997]: 'The Dissemination of Culture: A Model with Local Convergence and Global Polarization', *Journal of Conflict Resolution*, **41**, pp. 203–26.
- Barrett, J. A., Cochran, C. T., Huttegger, S. and Fujiwara, N. [2017]: 'Hybrid Learning in Signalling Games', *Journal of Experimental and Theoretical Artificial Intelligence*, **29**, pp. 1119–27.

- Barrett, J. and Zollman, K. J. S. [2009]: 'The Role of Forgetting in the Evolution and Learning of Language', *Journal of Experimental and Theoretical Artificial Intelligence*, **21**, pp. 293–309.
- Barton, N. H. [1979]: 'The Dynamics of Hybrid Zones', *Heredity*, **43**, pp. 685–97.
- Barton, N. H. and Hewitt, G. M. [1981]: 'Hybrid Zones and Speciation', in W. R. Atchley and D. S. Woodruff (eds), *Evolution and Speciation*, Cambridge: Cambridge University Press, pp. 109–45.
- Beggs, A. W. [2005]: 'On the Convergence of Reinforcement Learning', *Journal of Economic Theory*, **122**, pp. 1–36.
- Bomze, I. M. [1986]: 'Non-cooperative Two-Person Games in Biology: A Classification', *International Journal of Game Theory*, **15**, pp. 31–57.
- Erev, I. and Roth, A. E. [1998]: 'Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria', *American Economic Review*, **86**, pp. 848–81.
- Eshel, I. [2005]: 'Asymmetric Population Games and the Legacy of Maynard Smith: From Evolution to Game Theory and Back?', *Theoretical Population Biology*, **68**, pp. 11–18.
- Eshel, I. and Sansone, E. [2001]: 'Multiple Asymmetry and Concord Resolutions of a Conflict', *Journal of Theoretical Biology*, **213**, pp. 209–22.
- Fortini, S. and Petrone, S. [2012]: 'The Evolution of Respect for Property', *Statistics and Probability Letters*, **82**, pp. 1521–29.
- Goyal, S. and Janssen, M. [1996]: 'Can We Rationally Learn to Coordinate?', *Theory and Decision*, **40**, pp. 29–49.
- Hamilton, W. D. [1967]: 'Extraordinary Sex Ratios', *Science*, **156**, pp. 477–88.
- Hopkins, E. and Posch, M. [2005]: 'Attainability of Boundary Points under Reinforcement Learning', *Games and Economic Behavior*, **53**, pp. 110–25.
- Hoppe, F. M. [1984]: 'Pólya-Like Urns and the Ewens' Sampling Formula', *Journal of Mathematical Biology*, **20**, pp. 91–94.
- Hoppe, F. M. [1987]: 'The Sampling Theory of Neutral Alleles and an Urn Model in Population Genetics', *Journal of Mathematical Biology*, **25**, pp. 123–59.
- LaCroix, T. [2020]: 'On Salience and Signaling in Sender–Receiver Games: Partial Pooling, Learning, and Focal Points', *Synthese*, **82**, pp. 1725–47.
- Lewis, D. [1969]: *Convention: A Philosophical Study*, Cambridge, MA: Harvard University Press.
- Maynard Smith, J. [1982]: *Evolution and the Theory of Games*, Cambridge: Cambridge University Press.
- Maynard Smith, J. and Parker, G. A. [1976]: 'The Logic of Asymmetric Contests', *Animal Behaviour*, **24**, pp. 159–75.
- Maynard Smith, J. and Price, G. R. [1973]: 'The Logic of Animal Conflict', *Nature*, **246**, pp. 15–18.
- Mehta, J., Starmer, C. and Sugden, R. [1994a]: 'Focal Points in Pure Coordination Games: An Experimental Investigation', *Theory and Decision*, **36**, pp. 163–85.
- Mehta, J., Starmer, C. and Sugden, R. [1994b]: 'The Nature of Salience: An Experimental Investigation of Pure Coordination Games', *American Economic Review*, **84**, pp. 658–73.
- Mesterton-Gibbons, M. and Adams, E. S. [2003]: 'Landmarks in Territory Partitioning: A Strategically Stable Convention?', *American Naturalist*, **161**, pp. 685–97.

- Mohseni, A. [2019]: 'Stochastic Stability and Disagreements between Dynamics', *Philosophy of Science*, **86**, pp. 497–521.
- Morsky, B. and Akçay, E. [2019]: 'Evolution of Social Norms and Correlated Equilibria', *Proceedings of the National Academy of Sciences USA*, **116**, pp. 8834–39.
- Otto, S. P. and Whitlock, M. C. [1997]: 'The Probability of Fixation in Populations of Changing Size', *Genetics*, **146**, pp. 723–33.
- Pemantle, R. [2007]: 'A Survey of Random Processes with Reinforcement', *Probability Surveys*, **4**, pp. 1–79.
- Reinhard, S. [1980]: 'A Note on Evolutionary Stable Strategies in Asymmetric Animal Conflicts', *Journal of Theoretical Biology*, **84**, pp. 93–101.
- Roth, A. E. and Erev, I. [1995]: 'Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term', *Games and Economic Behavior*, **8**, pp. 164–212.
- Schelling, T. C. [2004]: *The Strategy of Conflict*, Cambridge, MA: Harvard University Press.
- Schlicht, E. [1998]: *On Custom in the Economy*, Oxford: Oxford University Press.
- Schreiber, S. J. [2001]: 'Urn Models, Replicator Processes, and Random Genetic Drift', *SIAM Journal on Applied Mathematics*, **61**, pp. 2148–67.
- Sherratt, T. N. and Mesterton-Gibbons, M. [2015]: 'The Evolution of Respect for Property', *Journal of Evolutionary Biology*, **28**, pp. 1185–202.
- Skyrms, B. [2010]: *Evolution, Learning, and Information*, Oxford: Oxford University Press.
- Sugden, R. [1995]: 'A Theory of Focal Points', *Economic Journal*, **105**, pp. 533–50.
- Sugden, R. [2004]: *The Economics of Rights, Co-operation, and Welfare*, Basingstoke: Springer.
- Sugden, R. [2011]: 'Salience, Inductive Reasoning, and the Emergence of Conventions', *Journal of Economic Behavior and Organization*, **79**, pp. 35–47.
- Vanderschraaf, P. [1995]: 'Convention as Correlated Equilibrium', *Erkenntnis*, **42**, pp. 65–87.
- Vickers, G. T. and Cannings, C. [1987]: 'On the Definition of an Evolutionarily Stable Strategy', *Journal of Theoretical Biology*, **129**, pp. 349–53.
- Wilson, D. S. [1975]: 'A Theory of Group Selection', *Proceedings of the National Academy of Sciences USA*, **72**, pp. 143–46.
- Wilson, D. S. and Sober, E. [1994]: 'Reintroducing Group Selection to the Human Behavioral Sciences', *Behavioral and Brain Sciences*, **17**, pp. 585–607.