Introduction to Game Theory

Part 3: More Advanced Concepts

3 Beyond normal-form games

Remark 3.1 (**A recapitulation**)
On the first day, we learnt the basics of game theory. We learnt how to formalize games, and how to solve them. And we learnt how to describe learning processes in games, and how to use simulations to explore these learning processes. Importantly, however, all the games we studied so far were rather simple. In particular, individuals only had to make one decision, they all moved simultaneously, and they had all relevant information when making a choice. While these are fair assumptions for some strategic decisions, many 'games' in real life are more complicated. For example, in some games, people make sequential decision (e.g., in this course, I first determined the rules how you get good grades; then you get a chance to adapt your strategy accordingly). Also, in some games incomplete information is an essential component (e.g., when bidding for a house, I do not know the seller's true reservation price). In this third part, we will discuss some of these more complex scenarios.

3.1 Sequential games

Remark 3.2 (**Sequential game**)
In the following, we say a game is sequential if it consists of several time steps (stages), and at any given stage only a single player is to make a decision.

Example 3.3 (**A sequential stag-hunt game**)
The setup. Suppose players interact in a stag-hunt game, with the same possible actions and the same payoffs as before,

	Stag	Hare
Stag	$10, 10$	$0, 4$
Hare	$4, 0$	$4, 4$

However, this time, the row-player makes her decision first, which is observed by the column-player. Then the column-player makes his own decision.

What difference does it make? The major difference to the previous setup is that now, the column-player can make his choice dependent on the row-player's decision. So for the row-player, the possible (pure) strategies are still

$$s_R \in \{\text{Stag}, \text{Hare}\}.$$

But for the column-player the possible strategies are now all pairs (a_S, a_H), where a_S is what the column player's reaction to Stag, and a_R is the column-player's reaction to Hare. Overall, there are four possibilities,

$$s_C \in \big\{(\text{Stag}, \text{Stag}),\ (\text{Stag}, \text{Hare}),\ (\text{Hare}, \text{Stag}),\ (\text{Hare}, \text{Hare})\big\};$$

One could refer to these four possibilities as *Always Stag*, *Do what the row-player did*, *Do the opposite of what the row-player did*, and *Always Hare*, respectively.

Analyzing sequential games as normal-form games. In principle, we could analyze this sequential game with our existing methods, by interpreting the game as a normal-form game. The respective payoff matrix would be

	(Stag,Stag)	(Stag,Hare)	(Hare,Stag)	(Hare,Hare)
Stag	$\underline{10, 10}$	$\underline{10, 10}$	$0, 4$	$0, 4$
Hare	$4, 0$	$4, 4$	$4, 0$	$\underline{4, 4}$

We would conclude that this game has three Nash equilibria (in pure strategies, underlined in the matrix). But are all of these Nash equilibria equally reasonable? [to be continued]

Remark 3.4 (**Game trees**)
Instead of using payoff matrices, an alternative way of representing games is to use game trees. This representation is particularly useful for sequential games, because it highlights the order of moves. Game trees consist of decision nodes connected by edges. Each decision node refers to a particular player who is to make a decision. The edges refer to the possible the actions the player can make at this point. In addition, there are terminal nodes that indicate that the game ends at that point. These terminal nodes also display the resulting payoffs to each player. For an illustration of the sequential stag-hunt game, see **Figure 1a**.

Example 3.5 (**A sequential stag-hunt game – continued**)
Revisiting the previous Nash equilibria. One can represent each possible strategy pair (s_R, s_C) by highlighting the respective decision the player would make at each decision node. In **Figure 1b–d**, we illustrate this approach for the three Nash equilibria we previously identified. In the equilibrium depicted in **Figure 1b**, it is somewhat unreasonable to assume that the column-player played Stag if the row-player went to play Hare (which she didn't). In the equilibrium depicted in **Figure 1d**, it is somewhat unreasonable to assume that the column-player played Hare if the row-player went to play Stag (which she didn't). So in both cases, the column-player is irrational "off the equilibrium path". In the solution depicted in **Figure 1d**, this is the only reason why the row-player does not deviate to Stag.

A conceptual conclusion. If players are rational, a reasonable requirement on a 'solution' would be:
1. Players at pre-terminal nodes make optimal decisions even in those nodes that are not reached
2. Their co-players take (1) into account and decide accordingly.

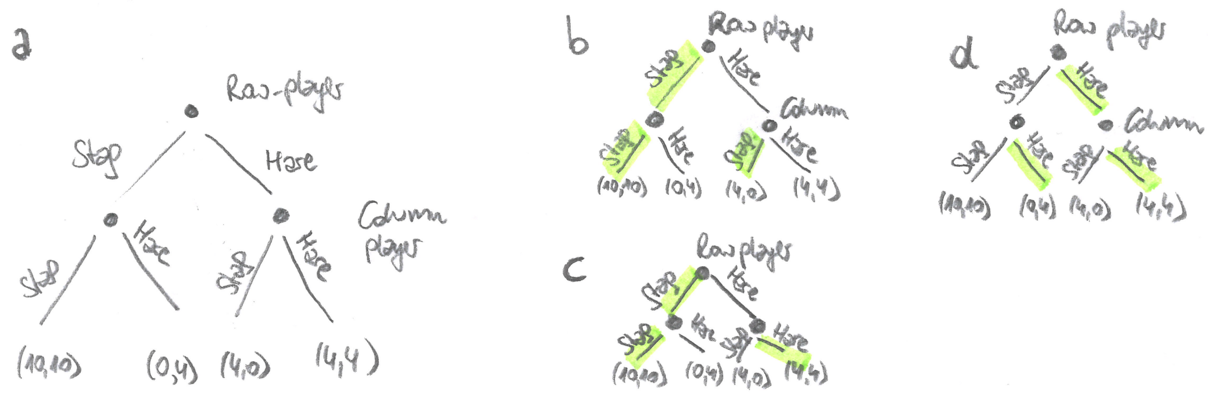

Figure 1: A game-tree analysis of the sequential hawk-dove game. a, We consider a version of the hawk-dove game where the row player moves first. **b–d,** An illustration of the three Nash equilibria of the normal-form representation of the game. The green edges represent the actions a player would take if the respective node is reached.

The unique solution that satisfies both of these requirements is the one depicted in **Figure 1c**. Here, the row-player uses Stag, and the column-player uses (Stag, Hare), i.e. do what the row-player did. What is interesting about this solution: now, among rational players, both players will choose Stag. This is different from the game where people had to choose simultaneously, where both players choosing Hare was also a Nash equilibrium [*How is this possible?*].

The previous idea of working backwards (from the bottom of the game tree to the top) can be generalized to game trees of arbitrary length (and for games with an arbitrary number of players).

Theorem 3.6 (**Backward induction.**)
For any sequential game with finitely many stages, one can find a Nash equilibrium with the following algorithm:

1. For each node in the final stage, determine the optimal choice
2. For each node in the last-to-final stage, determine the optimal choice, given play in the final stage is determined by step (1)
3. Roll back until the initial stage is reached.

This solution has the property that each player's action is optimal at every possible node. For *almost all* games (roughly: if you draw payoffs randomly), this solution is unique.

Group Exercise 3.7 (**Solving games by backward induction**)
Solve the following games by backward induction, and interpret the result (**Bonus:** Can you find other Nash equilibria?)

1. **Cooperation and punishment.** Consider a social interaction among two people, player 1 and player 2. In the first stage, player 1 can decide to help player 2. Helping here means to pay some small cost $c > 0$ to give a benefit $b > c$ to the co-player). In the second stage, player 2 can decide whether to punish

player 1. Punishing means to pay some small cost k to the reduce the co-player's payoff by p. Can you draw a game tree for this interaction? Can you solve the game by backward induction? Perhaps you could even use this game to discuss the issue of *(non-)credible threats?*

2. **The ultimatum game.** Suppose there are two players, a seller and a buyer. The seller has an item that is worth 100 EUR to the buyer, and that is worth nothing to the seller. Now they wish to bargain the price of the item. Suppose the rules are as follows: First, the seller announces some price $p \in \{1, 2, \ldots, 99\}$. Then the buyer either accepts the price or rejects the price. If the buyer accepts, players exchange the item, and the seller's payoff is p, and the buyer's payoff is $100-p$. If the buyer rejects, both players end up with a payoff of zero. What result would you expect with backward induction? Perhaps you can even use this result to discuss the technical term of having a *first-mover advantage*?

Remark 3.8 (**A critique of backward induction**)

1. Consider the example of a game depicted in **Figure 2a** among n players. Backward induction would predict that in this game, all players would always go right (choose R), to yield the payoff of 2 for everyone. However, now suppose player 1 thinks there is some small probability $\varepsilon > 0$ that any single co-player chooses down (D). Therefore, even in the best case, the chance to reach the payoff of 2 is $(1-\varepsilon)^n$. This is quite risky when the number n of players is large. Even if player 1 prefers to go right after these considerations, she might worry that player 2 engages in the same considerations and goes for the safe option. Backward induction requires long chains of "Player 1 knows that player 2 knows that etc". So for large n, backward induction is somewhat less compelling as a solution concept.

2. **Backward induction paradox.** Now consider the so-called *centipede game*, a game between two players depicted in **Figure 2b**. Here, backward induction suggests that in every decision node, players should move down (D). But is this solution compelling?
   Suppose you are player 2 and you unexpectedly find yourself in a situation where player 1 chose R_1 in the first round. Then backward induction suggests you should go down (D_2), for otherwise player 1 will play D_3. However, backward induction also says that player 1 should have never played R_1 in the first place. So it may be reasonable to choose R_2 after all. However, knowing this, player 1 may actually play R_1 on purpose.
   This so-called backward induction paradox arises because in the early stages of a sequential game, players are able to signal that they very assumptions of game theory (the usual kind of rationality) do not apply to them.
   [Some suggestions to fix this paradox: (1) There is some uncertainty in the payoffs of the game. Observing R_1 then just means that you misinterpreted player 1's payoff from D_1. (2) Sometimes deviations occur because there are rare and independent errors. So if player 2 observes R_1, he should assume this was just a one-time mistake and that player 2 would choose D_3 when given the chance.]

Remark 3.9 (**Experimental evidence for or against backward induction**)

- **Ultimatum game.** In experimental ultimatum games, a majority of sellers offers 40-50% of the surplus the buyer. Offers where the buyer gets less than 20% of the share are usually rejected. However, there

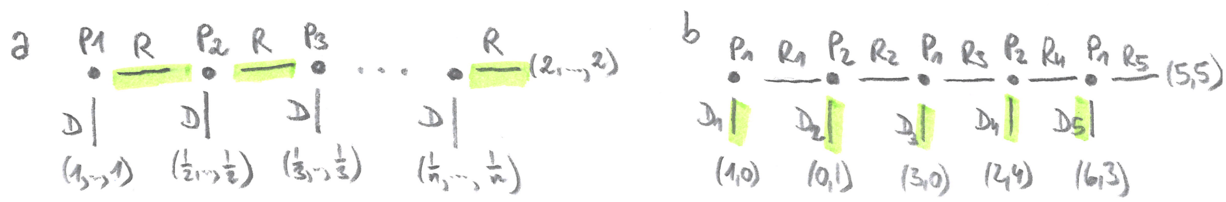

Figure 2: Two examples to discuss the plausibility of backward induction. The green lines represent the player's preferred actions at the respective node. **a,** The backward induction outcome becomes less compelling when the number of players grows large. **b,** In the centipede game, a backward induction paradox might arise: if player 2 is to make a decision, he might question player 1's rationality (which is the very baseline assumption of game theory).

is huge cross-cultural variation (Henrich et al., 2001) and outcomes depend on how much money is at stake (Andersen et al., 2011).

- **Centipede games.** Palacios-Huerta and Volij (2009) considered centipede games that involved chess players (i.e., a subject pool where you can assume a certain degree of rationality and foresight). At a first field experiment (conducted at various tournaments), 69% of games terminated at the very first node; this number increased to 100% if the first player was a grandmaster. In a second laboratory tournament, experimenters had a mixed subject pool of chess players and students. If two students interacted, 3% of games ended in the first round; if a student and a chess-player interact, this number changes to 30% (if the student moves first) and 37.5% (if the chess player moves first), respectively.

3.2 Repeated games

Remark 3.10 (**On repeated games**)
So far, we considered games in which players interact over multiple stages. However, we made the somewhat restrictive assumption that at each stage, only a single player is allowed to make a decision. Instead, one could consider so-called *multi-stage* games, where there also may be stages in which two or more players who decide simultaneously. One particularly important class of such multi-stage games are so-called *repeated games*. Repeated games are the type of strategic interaction that arise when you engage in the same normal-form game several times, with the same co-players.

Remark 3.11 (**The repeated prisoner's dilemma**)
Perhaps the most well-known repeated game is the repeated prisoner's dilemma. In the prisoner's dilemma, there are two players who can either cooperate or defect. Mutual cooperation gives a higher payoff than mutual defection, yet there is always some incentive to defect. A typical payoff matrix looks like this:

	Cooperate (C)	Defect (D)
Cooperate (C)	$3, 3$	$0, 5$
Defect (D)	$5, 0$	$1, 1$

As a normal-form game, this game has a unique Nash equilibrium, which is for both players to defect [*Which game we previously encountered has the same incentive structure as the prisoner's dilemma?*] However, researchers are also interested in the repeated version of this game (where players interact for 2 or more consecutive rounds). Now the players can use more complicated strategies. Instead of just choosing whether to cooperate or to defect (or how to randomize between them), they could follow complicated rules that tell them whether to cooperate in a given round [*Can you come up with some random examples of interesting strategies?*] So at least from the outset, exploring such repeated games seems to be more complex. There are two key motivations to do it anyway:

1. A theoretical one. For theorists, it is an interesting question to ask: Given we understand the equilibria of the one-shot game, would repetition ever make a difference?
2. A practical one. The repeated prisoner's dilemma seems to be a good model to explore the spontaneous emergence of reciprocal cooperation. For example, there is good evidence for reciprocal cooperation among humans (exchanging favors among friends), but also in several animal species, such as bats (Wilkinson, 1984) or Norway rats (Schweinfurth et al., 2019).

Remark 3.12 (**Does repetition make a difference?**)
Lacking access to bats or norway rats, let's tackle the first question above. Is there a reason to expect that rational individuals would behave any different in the repeated prisoner's dilemma, compared to the usual 'one-shot' game? The answer is: *It depends.*

- Finitely repeated prisoner's dilemma. Suppose the two players know beforehand for how many rounds they interact. Then the game can again be solved with a backward-induction argument. In the last round, players clearly have no incentive to cooperate, and hence they surely defect. But given that they certainly defect in the last round no matter what, they also have no incentive to cooperate in the second-to-last round either. Repeat this argument, and voilà, you observe that individuals should defect in every single round. We get the same outcome as in the one-shot game.

- Infinitely or indefinitely repeated prisoner's dilemma. Instead, suppose players don't know when the game ends (e.g., because there is always some positive probability to meet again). Now backward induction does not apply anymore, because there is no known last round where to start. And indeed, if the game is repeated for sufficiently many rounds, one can now construct Nash equilibria where players cooperate in every single round. To get some intuition, suppose after any round there is another one with 90% probability (or equivalently, that players interact for 10 rounds) Moreover, suppose the two players adopt the strategy `GRIM`. `GRIM` is the strategy that starts cooperating, and continues cooperating as long as the co-player cooperated every round so far. If the co-player ever deviates, `GRIM` would defect forever. Then if two `GRIM` players interact, they would cooperate in every round. As a result, each of them would get on average

$$10 \text{ [rounds in expectation]} \cdot 3 \text{ [mutual cooperation payoff]} \;=\; 30 \text{ [Total payoff]}$$

  To be a Nash equilibrium, no player must have an incentive to deviate. Suppose one player considered deviating to, say, the strategy `ALLD` that always defects. That deviation would pay in the very first round (getting 5 instead of 3). However, in all remaining rounds that player would end up with the mutual defection payoff of 1. In expectation, the resulting payoff is

$$5 + 9 \times 1 \;=\; 14 \text{ [Total payoff]}$$

Hence this deviation is not profitable (and neither is any other). That is, both players choosing `GRIM` is a Nash equilibrium. We conclude that the repeated game can lead to outcomes that are very different from the one-shot game.

Remark 3.13 (The folk theorem of repeated games)
If the repeated prisoner's dilemma now allows for more than one equilibrium, a natural question is: how many (qualitatively different) equilibria are there? Unfortunately, this number is vast – as the continuation probability approaches one (as the game is repeated for a very long time), there are infinitely many of them (e.g., Friedman, 1971). So how to predict anything for such repeated games?

Remark 3.14 (A famous tournament)
An alternative approach to make progress: Why not use tournament as a means to identify effective strategies for the repeated prisoner's dilemma? Axelrod and Hamilton (1981) asked fourteen scientists from different areas to submit strategies for a round-robin tournament. The scientists knew the payoff matrix (same as before), but they did not know the number of rounds each pairwise game is played (200 rounds). Some of the submitted strategies were rather sophisticated, trying to identify the co-player's type in the early rounds just to exploit the co-player later on. In spite of this, the most simple strategy (among the submitted ones) won, *Tit-for-Tat* (in the first round cooperate, then do whatever the co-player did in the previous round). However, that result is quite sensitive on the set of strategies that participated – so there is still a lot of research on optimal play in repeated games (for more on that, see for example Hilbe et al., 2018).

3.3 Games with incomplete information

Remark 3.15 (Motivation)
In the games so far, we made rather strong assumptions on what individuals know. Not only did they know their own available strategies and their own incentives to choose each strategy; they also had this kind of information for each co-player. Obviously, in practice there are many situations in which these assumptions do not hold. For example, in poker, individuals know their own cards, but not the cards of their opponents. When buying a house, I know my own valuation of the property, but I do not know the reservation price of the seller; nor do I know how many other potential buyers there are (*who are even the other players?*) In these example, players may try to use their co-player's past actions to update their beliefs on aspects they do not know. As one might expect, incorporating these kind of indirect inferences make game-theoretic models more complex. As a result, within this course, we will not be able to go into the details. Still, it is interesting to get a feeling for the resulting questions by looking at an example.

Example 3.16 (PhD applications)
Game setup. Let us consider a stylized example of how universities might choose their PhD students.

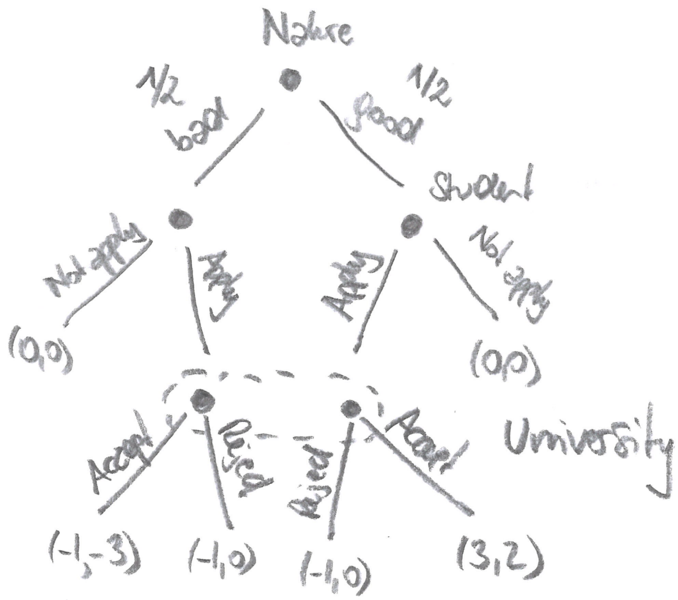

Figure 3: A game-tree representation of the PhD application game. The initial move by nature represents the uncertainty in the student's type – they can either be good or bad. The *information set* indicated by the dashed ellipse indicates that at each node in this ellipse, the university only knows that a student has applied, but not which type of student it is.

To this end, suppose students can either be extremely motivated and smart, or lazy and disinterested (for simplicity, we refer to these two students as 'good' or 'bad'). A student knows its own type (whereas the university only knows the general proportion of good students in the population). Based on their type, students decide whether to apply to the university. Applying to a university is costly, as it takes some time to prepare convincing application documents. However, if being admitted to the university, students may derive some positive benefit (a proper education) and so does the university (by gaining access to a strong researcher). Suppose the exact incentives are given by the game tree in **Figure 3**.

Players' strategies. As usual, strategies are rules that tell the player what to do, given the information they have. Students do know their type; so their strategies look like *Never apply*, *Apply if good*, *Apply if bad*, or *Always apply*. Now, the university in our example cannot learn anything about any applicant's type; they can only choose to *accept* any given applicant, or to *reject* the applicant.

On possible equilibria. Of course, from the viewpoint of the university, the best optimal outcome would be: Only good students apply, and those applicants are also accepted by the university. Is this an equilibrium? To answer that question, we need to ask whether anyone has an incentive to deviate.

- *Would bad students have an incentive to deviate?* According to this equilibrium bad students would not apply, and hence their payoff would be 0. If they applied instead, they would get accepted, leaving them with a payoff of -1. Hence, they would not want to deviate.

- *Would good students want to deviate?* In equilibrium, good students get accepted, and hence get a payoff of 3; by not applying they only get 0. Hence, again there is no profitable deviation.

- *Would the university want to deviate?* In equilibrium, the university accepts all applicants and those all happen to be good. Hence the university's payoff here is 2, whereas by not accepting applicants it would get 0. Again, no profitable deviation.

Therefore, this is an equilibrium. This equilibrium has the property that different types of students choose different actions. That is, in this equilbrium, even if the university does not know the students' types in advance, they can perfectly infer these types from the students' behavior. This kind of equilibrium is thus sometimes called a *separating equilibrium* or a *revealing equilibrium*. This is not the only kind of equilibrium; in fact, one can also show that this game allows for an equilibrium in which no student applies, and the university would reject if anyone applies. Obviously, in this equilibrium, students do not reveal their types; this is called a *pooling equilibrium*. These types of equilibria play a huge role in these so-called *signaling games*. This type of model can be used to explain, for example, why students go to universities in the first place (Spence, 1973), or why some people prefer to be seen as modest (Hoffman et al., 2018).

Example 3.17 (A fun example)
Using signaling games, one can also explain why scam/phishing emails are often so *bad* (i.e., they are so obviously fraudulent). Can you come up with an explanation?

Remark 3.18 (A summary)
In this third part of the course, you should have learned

- how to think about strategic interactions in which individuals make their decisions consecutively
- how researchers use repeated games to study reciprocal cooperation (more on that later!), and
- how to incorporate individuals' lack of knowledge into game-theoretic models.

Remark 3.19 (Further reading)
If you would like to learn more about game theory, there are some great books out there. For example, if you would like to have a similarly low-level introduction to game theory, you could read Binmore (2007). For mathematical details, there is the classical book by Fudenberg and Tirole (1998) among many others. Finally, for people interested in evolutionary game theory, a good source is the book of Nowak (2006), or if you want to see even more sophisticated mathematics, the book by Hofbauer and Sigmund (1998).

References

Andersen, S., Ertaç, S., Gneezy, U., Hoffman, M., and List, J. A. (2011). Stakes matter in ultimatum games. *American Economic Review*, 101:3427–3439.

Axelrod, R. and Hamilton, W. D. (1981). The evolution of cooperation. *Science*, 211:1390–1396.

Binmore, K. (2007). *Game Theory: A very short introduction*. Oxford University Press.

Friedman, J. (1971). A non-cooperative equilibrium for supergames. *Review of Economic Studies*, 38:1–12.

Fudenberg, D. and Tirole, J. (1998). *Game Theory*. MIT Press, Cambridge, 6th edition.

Henrich, J., Boyd, R., Bowles, S., Camerer, C., Gintis, H., McElreath, R., and Fehr, E. (2001). In search of homo economicus: Experiments in 15 small-scale societies. *American Economic Review*, 91:73–79.

Hilbe, C., Chatterjee, K., and Nowak, M. A. (2018). Partners and rivals in direct reciprocity. *Nature Human Behaviour*, 2(7):469–477.

Hofbauer, J. and Sigmund, K. (1998). *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge, UK.

Hoffman, M., Hilbe, C., and Nowak, M. A. (2018). The signal-burying game can explain why we obscure positive traits and good deeds. *Nature Human Behaviour*, 2:397–404.

Nowak, M. A. (2006). *Evolutionary dynamics*. Harvard University Press, Cambridge MA.

Palacios-Huerta, I. and Volij, O. (2009). Field centipedes. *American Economic Review*, 99:1619–1635.

Schweinfurth, M. K., Aeschbacher, J., Santi, M., and Taborsky, M. (2019). Male norway rats cooperate according to direct but not generalized reciprocity rules. *Animal Behaviour*, 152:93–101.

Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics*, 87(3):355–374.

Wilkinson, G. S. (1984). Reciprocal food-sharing in the vampire bat. *Nature*, 308:181–184.