Game Theory

2.6 Application II: Convention

George Mason University, Spring 2018

Lewis' *Convention* [1969]

It is often said that language is conventional. But what does this mean exactly? It is not the case that all of our linguistic conventions could have been agreed upon by a board of syndics. After all, they would have had to speak a rudimentary language to hash out the terms of their agreement.

Lewis' *Convention* [1969]

It is often said that language is conventional. But what does this mean exactly? It is not the case that all of our linguistic conventions could have been agreed upon by a board of syndics. After all, they would have had to speak a rudimentary language to hash out the terms of their agreement.

Lewis' aim is to analyze convention in its full generality, including tacit convention not created by explicit agreement.

I. Coordination and Convention

- Coordination Problems
- Solving Coordination Problems
- Convention (First Pass)

Ex. (1) Meeting (cf. Bach or Stravinsky?).

"Suppose you and I both want to meet each other. We will meet if and only if we go to the same place. It matters little to either of us where (within limits) he goes if he meets the other there; and it matters little to either of us where he goes if he fails to meet the other there. We must each choose where to go. The best place for me to go is the place where you will go, so I try to figure out where you will go and to go there myself. You do the same. Each chooses according to his expectation of the others choice. If either succeeds, so does the other; the outcome is one we both desired."

Ex. (2) Telephone.

"Suppose you and I are talking on the telephone and we are unexpectedly cut off after three minutes. We both want the connection restored immediately, which it will be if and only if one of us calls back while the other waits. It matters little to either of us whether he is the one to call back or the one to wait. We must each choose whether to call back, each according to his expectation of the other's choice, in order to call back if and only if the other waits."

Ex. (3) Rowing (Hume).

"Suppose you and I are rowing a boat together. If we row in rhythm, the boat goes smoothly forward; otherwise the boat goes slowly and erratically, we waste effort, and we risk hitting things. We are always choosing whether to row faster or slower; it matters little to either of us at what rate we row, provided we row in rhythm. So each is constantly adjusting his rate to match the rate he expects the other to maintain."

Ex. (4) Driving.

"Suppose several of us are driving on the same winding two-lane roads. It matters little to anyone whether he drives in the left or the right lane, provided the others do likewise. But if some drive in the left lane and some in the right, everyone is in danger of collision. So each must choose whether to drive in the left lane or in the right, according to his expectations about the others: to drive in the left lane if most or all of the others do, to drive in the right lane if most or all of the others do (and to drive where he pleases if the others are more or less equally divided)."

Ex. (8) Stag Hunt (Rousseau).

"Suppose we are in a wilderness without food. Separately we can catch rabbits and eat badly. Together we can catch stags and eat well. But if even one of us deserts the stag hunt to catch a rabbit, the stag will get away; so the other stag hunters will not eat unless they desert too. Each must choose whether to stay with the stag hunt or desert according to his expectations about the others, staying if and only if no one else will desert."

Ex. (11) Language.

“Suppose that with practice we could adopt any language in some wide range. It matters comparatively little to anyone (in the long run) what language he adopts, so long as he and those around him adopt the same language and can communicate easily. Each must choose what language to adopt according to his expectations about his neighbors' language: English among English speakers, Welsh among Welsh speakers, Esperanto among Esperanto speakers, and so on.”

What do these *coordination problems* have in common? What are their important features?

Def 2.6.1. An n-player game of *pure conflict* $\mathcal{G}$, or *zero sum* game, is one where for any outcome $o \in \mathcal{O}$, $u_1(o) + ... + u_n(o) = 0$ (under some linear rescaling). That is, in any square, all players utilities sum to zero.

Def 2.6.1. An *n*-player game of *pure conflict* $\mathcal{G}$, or *zero sum* game, is one where for any outcome $o \in \mathcal{O}$, $u_1(o) + ... + u_n(o) = 0$ (under some linear rescaling). That is, in any square, all players utilities sum to zero.

Def 2.6.2. An *n*-player game of *pure coordination* $\mathcal{G}$ is one where for any outcome $o \in \mathcal{O}$, $u_1(o) = ... = u_n(o)$ (under some linear rescaling). That is, in any square, the utilities of all players are equal to one another.

Def 2.6.1. An *n*-player game of *pure conflict* $\mathcal{G}$, or *zero sum* game, is one where for any outcome $o \in \mathcal{O}$, $u_1(o) + ... + u_n(o) = 0$ (under some linear rescaling). That is, in any square, all players utilities sum to zero.

Def 2.6.2. An *n*-player game of *pure coordination* $\mathcal{G}$ is one where for any outcome $o \in \mathcal{O}$, $u_1(o) = ... = u_n(o)$ (under some linear rescaling). That is, in any square, the utilities of all players are equal to one another.

There is a spectrum between games of pure conflict and games of pure coordination. Coordination games are closer to the latter end of this spectrum.

Def 2.6.1. An n-player game of *pure conflict* $\mathcal{G}$, or *zero sum* game, is one where for any outcome $o \in \mathcal{O}$, $u_1(o) + ... + u_n(o) = 0$ (under some linear rescaling). That is, in any square, all players utilities sum to zero.

Def 2.6.2. An n-player game of *pure coordination* $\mathcal{G}$ is one where for any outcome $o \in \mathcal{O}$, $u_1(o) = ... = u_n(o)$ (under some linear rescaling). That is, in any square, the utilities of all players are equal to one another.

There is a spectrum between games of pure conflict and games of pure coordination. Coordination games are closer to the latter end of this spectrum.

The players' interests rise and fall together.

Ex. Meeting.

	London	Tokyo	Anchorage
London	3,2	1,0	1,1
Tokyo	1,0	2,2	0,1
Anchorage	2,1	0,1	1,1

Ex. Meeting.

	London	Tokyo	Anchorage
London	3,2	1,0	1,1
Tokyo	1,0	2,2	0,1
Anchorage	2,1	0,1	1,1

The maximum difference in a single cell is 1.
The maximum difference across cells is 3.
In such cases, Lewis says, the "coincidence of interests predominates."

Ex. Telephone.

	call back	wait
call back	-1,-1	1,1
wait	1,1	-1,-1

Ex. Telephone.

	call back	wait
call back	-1,-1	1,1
wait	1,1	-1,-1

The maximum difference in a single cell is 0.
The maximum difference across cells is 2.

Ex. Telephone.

	call back	wait
call back	-1,-1	1,1
wait	1,1	-1,-1

The maximum difference in a single cell is 0.
The maximum difference across cells is 2.

This is a game of pure coordination.

Ex. Rowing.

	row fast	row slow
row fast	2,2	-2,-1
row slow	-1,-2	1,1

Ex. Rowing.

	row fast	row slow
row fast	2,2	-2,-1
row slow	-1,-2	1,1

The maximum difference in a single cell is 1.
The maximum difference across cells is 4.

Ex. Stag Hunt.

	hunt stag	hunt hare
hunt stag	2,2	0,1
hunt hare	1,0	1,1

Ex. Stag Hunt.

	hunt stag	hunt hare
hunt stag	2,2	0,1
hunt hare	1,0	1,1

The maximum difference in a single cell is 1.
The maximum difference across cells is 2.

Def 2.6.3. A *coordination equilibrium* of $\mathcal{G}$ is an action profile $a^* \in \times_{i \in \mathcal{N}} \mathcal{A}_i$ such that for players $j, k \in \mathcal{N}$, the following condition holds:

$u_k(g(\langle a_j^*, a_{-j}^* \rangle)) \geq u_k(g(\langle a_j, a_{-j}^* \rangle))$ for each $a_j \in \mathcal{A}_j$.

Def 2.6.3. A *coordination equilibrium* of $\mathcal{G}$ is an action profile $a^* \in \times_{i \in \mathcal{N}} \mathcal{A}_i$ such that for players $j, k \in \mathcal{N}$, the following condition holds:

$u_k(g(\langle a_j^*, a_{-j}^* \rangle)) \geq u_k(g(\langle a_j, a_{-j}^* \rangle))$ for each $a_j \in \mathcal{A}_j$.

A coordination equilibrium is a combination in which no one would have been better off had *any one* agent alone had acted otherwise, either himself or someone else.

Def 2.6.3. A *coordination equilibrium* of $\mathcal{G}$ is an action profile $a^* \in \times_{i \in \mathcal{N}} \mathcal{A}_i$ such that for players $j, k \in \mathcal{N}$, the following condition holds:

$u_k(g(\langle a_j^*, a_{-j}^* \rangle)) \geq u_k(g(\langle a_j, a_{-j}^* \rangle))$ for each $a_j \in \mathcal{A}_j$.

A coordination equilibrium is a combination in which no one would have been better off had *any one* agent alone had acted otherwise, either himself or someone else.

Coordination equilibria are clearly Nash equilibria because of the cases where $j = k$.

Def 2.6.3. A *coordination equilibrium* of $\mathcal{G}$ is an action profile $a^* \in \times_{i \in \mathcal{N}} \mathcal{A}_i$ such that for players $j, k \in \mathcal{N}$, the following condition holds:

$u_k(g(\langle a_j^*, a_{-j}^* \rangle)) \geq u_k(g(\langle a_j, a_{-j}^* \rangle))$ for each $a_j \in \mathcal{A}_j$.

A coordination equilibrium is a combination in which no one would have been better off had *any one* agent alone had acted otherwise, either himself or someone else.

Coordination equilibria are clearly Nash equilibria because of the cases where $j = k$.

However, the converse does not hold.

Ex. Meeting.

	London	Tokyo	Anchorage
London	3,2	1,0	1,1
Tokyo	1,0	2,2	0,1
Anchorage	2,1	0,1	1,2

Ex. Meeting.

	London	Tokyo	Anchorage
London	3,2	1,0	1,1
Tokyo	1,0	2,2	0,1
Anchorage	2,1	0,1	1,2

$\langle$London, London$\rangle$ is a coordinated equilibrium.

Ex. Meeting.

	London	Tokyo	Anchorage
London	3,2	1,0	1,1
Tokyo	1,0	2,2	0,1
Anchorage	2,1	0,1	1,2

⟨Tokyo, Tokyo⟩ is a coordinated equilibrium.

Ex. Meeting.

	London	Tokyo	Anchorage
London	3,2	1,0	1,1
Tokyo	1,0	2,2	0,1
Anchorage	2,1	0,1	1,2

$\langle$Anchorage, Anchorage$\rangle$ is a Nash equilibrium
but not a coordinated equilibrium.

Ex. Telephone.

	call back	wait
call back	-1,-1	1,1
wait	1,1	-1,-1

Ex. Telephone.

	call back	wait
call back	-1,-1	1,1
wait	1,1	-1,-1

$\langle$call back, wait$\rangle$ and $\langle$wait, call back$\rangle$ are coordinated equilibria.

Ex. Rowing.

	row fast	row slow
row fast	2,2	-2,-1
row slow	-1,-2	1,1

Ex. Rowing.

	row fast	row slow
row fast	2,2	-2,-1
row slow	-1,-2	1,1

$\langle$row fast, row fast$\rangle$ and $\langle$row slow, row slow$\rangle$ are coordinated equilibria.

Ex. Stag Hunt.

	hunt stag	hunt hare
hunt stag	2,2	0,1
hunt hare	1,0	1,1

Ex. Stag Hunt.

	hunt stag	hunt hare
hunt stag	2,2	0,1
hunt hare	1,0	1,1

$\langle$hunt stag, hunt stag$\rangle$ and $\langle$hunt hare, hunt hare$\rangle$ are coordinated equilibria.

Any game of pure coordination will have at least one coordinated equilibrium.

Any game of pure coordination will have at least one coordinated equilibrium.

For a problem to be a *coordination* problem, it must have two or more different coordinated equilibria.

Any game of pure coordination will have at least one coordinated equilibrium.

For a problem to be a *coordination* problem, it must have two or more different coordinated equilibria.

But this requirement is not quite strong enough.

Any game of pure coordination will have at least one coordinated equilibrium.

For a problem to be a *coordination* problem, it must have two or more different coordinated equilibria.

But this requirement is not quite strong enough.

	a_1	a_2
a_1	1,1	1,1
a_2	0,0	0,0

$\langle a_1, a_1 \rangle$ and $\langle a_1, a_2 \rangle$ are coordinated equilibria but there is no real problem here: player 1 can ensure coordination by choosing act a_1.

Def 2.6.4. A *proper equilibrium* of $\mathcal{G}$ is an action profile $a^* \in \times_{i \in \mathcal{N}} \mathcal{A}_i$ such that for $j \in \mathcal{N}$, the following condition holds:

A proper equilibrium is one that each agent likes *better than* any other outcome he could have reached given the others choices.

$u_j(g(\langle a_j^*, a_{-j}^* \rangle)) > u_j(g(\langle a_j, a_{-j}^* \rangle))$ for each $a_j \in \mathcal{A}_j$.

Def 2.6.4. A *proper equilibrium* of $\mathcal{G}$ is an action profile $a^* \in \times_{i \in \mathcal{N}} \mathcal{A}_i$ such that for $j \in \mathcal{N}$, the following condition holds:

A proper equilibrium is one that each agent likes *better than* any other outcome he could have reached given the others choices.

$u_j(g(\langle a_j^*, a_{-j}^* \rangle)) > u_j(g(\langle a_j, a_{-j}^* \rangle))$ for each $a_j \in \mathcal{A}_j$.

	a_1	a_2
a_1	1,1	1,1
a_2	0,0	0,0

Neither $\langle a_1, a_1 \rangle$ nor $\langle a_1, a_2 \rangle$ is a proper equilibrium.

Def 2.6.4. A *proper equilibrium* of $\mathcal{G}$ is an action profile $a^* \in \times_{i \in \mathcal{N}} \mathcal{A}_i$ such that for $j \in \mathcal{N}$, the following condition holds:

A proper equilibrium is one that each agent likes *better than* any other outcome he could have reached given the others choices.

$u_j(g(\langle a_j^*, a_{-j}^* \rangle)) > u_j(g(\langle a_j, a_{-j}^* \rangle))$ for each $a_j \in \mathcal{A}_j$.

	a_1	a_2
a_1	1,1	1,1
a_2	0,0	0,0

Neither $\langle a_1, a_1 \rangle$ nor $\langle a_1, a_2 \rangle$ is a proper equilibrium.

A coordination problem will have two or more proper coordinated equilibria.

Isn't this a genuine coordination problem?

	a_1	a_2	a_3
a_1	1,1	1,1	0,0
a_2	1,1	1,1	0,0
a_3	0,0	0,0	1,1

Isn't this a genuine coordination problem?

	a_1	a_2	a_3
a_1	1,1	1,1	0,0
a_2	1,1	1,1	0,0
a_3	0,0	0,0	1,1

Intuitively, yes. But this problem can be reformulated as follows:

	a_1 or a_2	a_3
a_1 or a_2	1,1	0,0
a_3	0,0	1,1

Def 2.6.5. A *coordination problem* is a situation of "interdependent decision by two or more agents in which coincidence of interest predominates and in which there are two or more proper coordinated equilibria."

Agents might succeed in coordinating by a stroke of good luck.

Agents might succeed in coordinating by a stroke of good luck.

However, they are more likely to reach a coordinated equilibrium “through the agency of a system of suitably concordant mutual expectations.”

Agents might succeed in coordinating by a stroke of good luck.

However, they are more likely to reach a coordinated equilibrium "through the agency of a system of suitably concordant mutual expectations."

"In general, each may do his part of one of the possible coordination equilibria because he expects the others to do theirs, thereby reaching that equilibrium."

Agents might succeed in coordinating by a stroke of good luck.

However, they are more likely to reach a coordinated equilibrium "through the agency of a system of suitably concordant mutual expectations."

"In general, each may do his part of one of the possible coordination equilibria because he expects the others to do theirs, thereby reaching that equilibrium."

"He has a decisive reason to do his own part if he is *sufficiently* confident in his expectation that the others will do theirs. The degree of confidence which is sufficient depends on all his payoffs and sometimes on the comparative probabilities he assigns to the different *ways* the others might not all do their parts, in case not all of them do."

	a_1	a_2	a_3
a_1	1,1	0,0	-8,-8
a_2	0,0	1,1	9,9

	a_1	a_2	a_3
a_1	1,1	0,0	-8,-8
a_2	0,0	1,1	9,9

Suppose that player 1 knows that if player 2 does not do her part in reaching the coordinated equilibrium $\langle a_1, a_1 \rangle$ by choosing a_1, then she will choose a_2.

	a_1	a_2	a_3
a_1	1,1	0,0	-8,-8
a_2	0,0	1,1	9,9

Suppose that player 1 knows that if player 2 does not do her part in reaching the coordinated equilibrium $\langle a_1, a_1 \rangle$ by choosing a_1, then she will choose a_2.

$EU_1(a_1) = 1 \times Cr_1(\text{player 2 performs } a_1) + 0 \times Cr_1(\text{player 2 performs } a_2).$

$EU_1(a_2) = 0 \times Cr_1(\text{player 2 performs } a_1) + 1 \times Cr_1(\text{player 2 performs } a_2).$

	a_1	a_2	a_3
a_1	1,1	0,0	-8,-8
a_2	0,0	1,1	9,9

Suppose that player 1 knows that if player 2 does not do her part in reaching the coordinated equilibrium $\langle a_1, a_1 \rangle$ by choosing a_1, then she will choose a_2.

$EU_1(a_1) = 1 \times Cr_1(\text{player 2 performs } a_1) + 0 \times Cr_1(\text{player 2 performs } a_2)$.

$EU_1(a_2) = 0 \times Cr_1(\text{player 2 performs } a_1) + 1 \times Cr_1(\text{player 2 performs } a_2)$.

Assuming $Cr_1(\text{player 2 performs } a_1) + Cr_1(\text{player 2 performs } a_2) = 1$,

$EU_1(a_1) > EU_1(a_2)$ iff $Cr_1(\text{player 2 performs } a_1) > 0.5$.

	a_1	a_2	a_3
a_1	1,1	0,0	-8,-8
a_2	0,0	1,1	9,9

Suppose that player 1 knows that if player 2 does not do her part in reaching the coordinated equilibrium $\langle a_1, a_1 \rangle$ by choosing a_1, then she will choose a_2.

$EU_1(a_1) = 1 \times Cr_1(\text{player 2 performs } a_1) + 0 \times Cr_1(\text{player 2 performs } a_2)$.

$EU_1(a_2) = 0 \times Cr_1(\text{player 2 performs } a_1) + 1 \times Cr_1(\text{player 2 performs } a_2)$.

Assuming $Cr_1(\text{player 2 performs } a_1) + Cr_1(\text{player 2 performs } a_2) = 1$,

$EU_1(a_1) > EU_1(a_2)$ iff $Cr_1(\text{player 2 performs } a_1) > 0.5$.

Thus, player 1 will do her part in reaching the coordinated equilibrium $\langle a_1, a_1 \rangle$ iff $Cr_1(\text{player 2 performs } a_1) > 0.5$.

	a_1	a_2	a_3
a_1	1,1	0,0	-8,-8
a_2	0,0	1,1	9,9

Suppose that player 1 knows that if player 2 does not do her part in reaching the coordinated equilibrium $\langle a_1, a_1 \rangle$ by choosing a_1, then she is just as likely to choose a_2 as a_3.

	a_1	a_2	a_3
a_1	1,1	0,0	-8,-8
a_2	0,0	1,1	9,9

Suppose that player 1 knows that if player 2 does not do her part in reaching the coordinated equilibrium $\langle a_1, a_1 \rangle$ by choosing a_1, then she is just as likely to choose a_2 as a_3.

$EU_1(a_1) = 1 \times Cr_1(\text{player 2 performs } a_1) - 8 \times Cr_1(\text{player 2 performs } a_3)$.

$EU_1(a_2) = 1 \times Cr_1(\text{player 2 performs } a_2) + 9 \times Cr_1(\text{player 2 performs } a_3)$.

	a_1	a_2	a_3
a_1	1,1	0,0	-8,-8
a_2	0,0	1,1	9,9

Suppose that player 1 knows that if player 2 does not do her part in reaching the coordinated equilibrium $\langle a_1, a_1 \rangle$ by choosing a_1, then she is just as likely to choose a_2 as a_3.

$EU_1(a_1) = 1 \times Cr_1(\text{player 2 performs } a_1) - 8 \times Cr_1(\text{player 2 performs } a_3)$.

$EU_1(a_2) = 1 \times Cr_1(\text{player 2 performs } a_2) + 9 \times Cr_1(\text{player 2 performs } a_3)$.

Since $Cr_1(\text{player 2 performs } a_2) = \frac{1 - Cr_1(\text{player 2 performs } a_1)}{2}$ and $Cr_1(\text{player 2 performs } a_3) = \frac{1 - Cr_1(\text{player 2 performs } a_1)}{2}$,

$EU_1(a_1) > EU_1(a_2)$ iff $Cr_1(\text{player 2 performs } a_1) > 0.9$.

	a_1	a_2	a_3
a_1	1,1	0,0	-8,-8
a_2	0,0	1,1	9,9

Suppose that player 1 knows that if player 2 does not do her part in reaching the coordinated equilibrium $\langle a_1, a_1 \rangle$ by choosing a_1, then she is just as likely to choose a_2 as a_3.

$EU_1(a_1) = 1 \times Cr_1(\text{player 2 performs } a_1) - 8 \times Cr_1(\text{player 2 performs } a_3)$.

$EU_1(a_2) = 1 \times Cr_1(\text{player 2 performs } a_2) + 9 \times Cr_1(\text{player 2 performs } a_3)$.

Since $Cr_1(\text{player 2 performs } a_2) = \frac{1 - Cr_1(\text{player 2 performs } a_1)}{2}$ and $Cr_1(\text{player 2 performs } a_3) = \frac{1 - Cr_1(\text{player 2 performs } a_1)}{2}$,

$EU_1(a_1) > EU_1(a_2)$ iff $Cr_1(\text{player 2 performs } a_1) > 0.9$.

Thus, player 1 will do her part in reaching the coordinated equilibrium $\langle a_1, a_1 \rangle$ iff $Cr_1(\text{player 2 performs } a_1) > 0.9$.

In some cases, it might be near impossible to coordinate given the high level of mutual confidence required. Example: The millionaire giving away his fortune to 1000 people.

In some cases, it might be near impossible to coordinate given the high level of mutual confidence required. Example: The millionaire giving away his fortune to 1000 people.

But we can often acquire the mutually concordant expectations required for coordination.

In some cases, it might be near impossible to coordinate given the high level of mutual confidence required. Example: The millionaire giving away his fortune to 1000 people.

But we can often acquire the mutually concordant expectations required for coordination.

One way to do this is to put ourselves in each other's shoes.

Ex. Meeting.

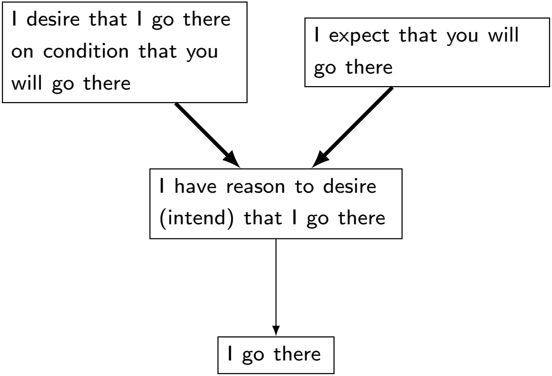
I desire that I go there on condition that you will go there
I expect that you will go there
I have reason to desire (intend) that I go there
I go there

Ex. Meeting.

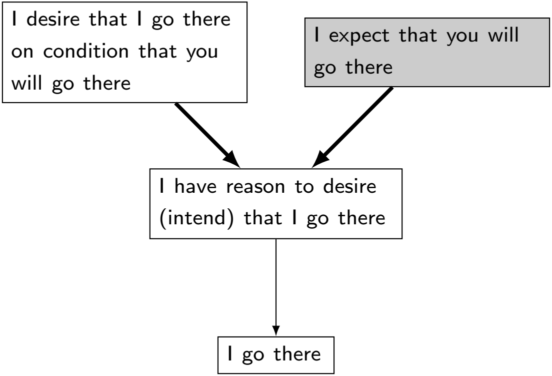
I desire that I go there on condition that you will go there
I expect that you will go there
I have reason to desire (intend) that I go there
I go there

Ex. Meeting.

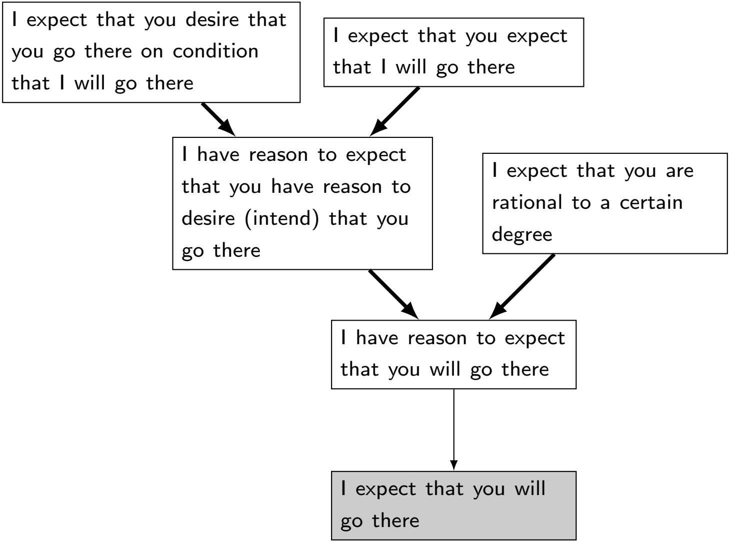
I expect that you desire that you go there on condition that I will go there
I expect that you expect that I will go there
I have reason to expect that you have reason to desire (intend) that you go there
I expect that you are rational to a certain degree
I have reason to expect that you will go there
I expect that you will go there

“Note that replication is *not* an interaction back and forth between people. It is a process in which *one* person works out the consequences of his beliefs about the world—a world he believes to include other people who are working out the consequences of their beliefs, including their belief in other people who...By our interaction in the world we acquire various high-order expectations that can serve us as premises. In our subsequent reasoning we are windowless monads doing our best to mirror each other, mirror each other mirroring each other, and so on.”

“Note that replication is *not* an interaction back and forth between people. It is a process in which *one* person works out the consequences of his beliefs about the world—a world he believes to include other people who are working out the consequences of their beliefs, including their belief in other people who...By our interaction in the world we acquire various high-order expectations that can serve us as premises. In our subsequent reasoning we are windowless monads doing our best to mirror each other, mirror each other mirroring each other, and so on.”

We will not generally solve a coordination problem by reasoning from 100th-order expectations. Nevertheless, higher-order expectations provide reasons to do one’s part.

“Note that replication is *not* an interaction back and forth between people. It is a process in which *one* person works out the consequences of his beliefs about the world—a world he believes to include other people who are working out the consequences of their beliefs, including their belief in other people who...By our interaction in the world we acquire various high-order expectations that can serve us as premises. In our subsequent reasoning we are windowless monads doing our best to mirror each other, mirror each other mirroring each other, and so on.”

We will not generally solve a coordination problem by reasoning from 100th-order expectations. Nevertheless, higher-order expectations provide reasons to do one’s part.

“The more orders, the better.”

One way to engender the concordant mutual expectations required to solve a coordination problem is through explicit agreement—a promise, declaration of intention, and so forth.

One way to engender the concordant mutual expectations required to solve a coordination problem is through explicit agreement—a promise, declaration of intention, and so forth.

If one promises to coordinate, one has a *second* independent incentive to coordinate. Indeed, strong promises might change the payoffs to such a degree that we no longer have a coordination problem.

One way to engender the concordant mutual expectations required to solve a coordination problem is through explicit agreement—a promise, declaration of intention, and so forth.

If one promises to coordinate, one has a *second* independent incentive to coordinate. Indeed, strong promises might change the payoffs to such a degree that we no longer have a coordination problem.

Ex. Stag Hunt.

	hunt stag	hunt hare
hunt stag	2,2	0,1
hunt hare	1,0	1,1

If breaking a promise lowers one's utility by 2, then the only Nash equilibrium is ⟨hunt stag, hunt stag⟩.

One way to engender the concordant mutual expectations required to solve a coordination problem is through explicit agreement—a promise, declaration of intention, and so forth.

If one promises to coordinate, one has a *second* independent incentive to coordinate. Indeed, strong promises might change the payoffs to such a degree that we no longer have a coordination problem.

Ex. Stag Hunt.

	hunt stag	hunt hare
hunt stag	2,2	0,-1
hunt hare	-1,0	-1,-1

If breaking a promise lowers one's utility by 2, then the only Nash equilibrium is $\langle$hunt stag, hunt stag$\rangle$.

However, explicit agreement is not the only source of concordant expectations.

However, explicit agreement is not the only source of concordant expectations.

"Explicit agreement is an especially good and common means to coordination—so much so that we are tempted to speak of coordination otherwise produced as *tacit* agreement. But agreement (literally understood) is not the only source of concordant expectations to help us solve our coordination problems. We do without agreement by choice if we find ourselves already satisfied with the content and strength of our mutual expectations. We do without it by necessity if we have no way to communicate, or if we can communicate only at a cost that outweighs our improved chance of coordination (say, if we are conspirators being shadowed)."

Experiments by Schelling reveal that we often do well at solving novel coordination problems without communicating. Subjects try to reach a coordination equilibrium that is *salient* in some respect (salience, rather than goodness, is what matters).

Experiments by Schelling reveal that we often do well at solving novel coordination problems without communicating. Subjects try to reach a coordination equilibrium that is *salient* in some respect (salience, rather than goodness, is what matters).

One source of salience is *precedent*. If players have already faced the coordination problem before, or a similar coordination problem before, then one equilibrium might be unique in a preeminently conspicuous respect because it, or an analogous equilibrium, was reached in the previous problem.

Experiments by Schelling reveal that we often do well at solving novel coordination problems without communicating. Subjects try to reach a coordination equilibrium that is *salient* in some respect (salience, rather than goodness, is what matters).

One source of salience is *precedent*. If players have already faced the coordination problem before, or a similar coordination problem before, then one equilibrium might be unique in a preeminently conspicuous respect because it, or an analogous equilibrium, was reached in the previous problem.

A fictive precedent can also do the trick. Example: Fabricated story about meeting on Charles Street.

There might be many precedents to follow. Players' actions might conform to a noticeable *regularity*.

There might be many precedents to follow. Players' actions might conform to a noticeable *regularity*.

We can reach coordination equilibria in new coordination problems by all continuing to conform to this same regularity.

There might be many precedents to follow. Players' actions might conform to a noticeable *regularity*.

We can reach coordination equilibria in new coordination problems by all continuing to conform to this same regularity.

It does not matter why coordination was achieved at particular equilibria in the past.

There might be many precedents to follow. Players' actions might conform to a noticeable *regularity*.

We can reach coordination equilibria in new coordination problems by all continuing to conform to this same regularity.

It does not matter why coordination was achieved at particular equilibria in the past.

If there is a regularity in action, each of us needn't be acquainted with the exact same past coordination problems in order to coordinate in future problems. Also, our acquaintance with a precedent needn't be detailed.

"Coordination by precedent, at its simplest, is this: achievement of coordination by means of shared acquaintance with the achievement of coordination in a single past case exactly like our present coordination problem. By removing inessential restrictions, we have come to this: achievement of coordination by means of shared acquaintance with a *regularity* governing the achievement of coordination in a class of past cases which bear some conspicuous analogy to one another and to our present coordination problem. Our acquaintance with this regularity comes from our experience with some of its instances, not necessarily the same ones for everybody."

“Each new action in conformity to the regularity adds to our experience of general conformity. Our experience of general conformity in the past leads us, by force of precedent, to expect a like conformity in the future. And our expectation of future conformity is a reason to go on conforming, since to conform if others do is to achieve a coordination equilibrium and to satisfy one’s own preferences. And so it goes—we’re here because we’re here because we’re here because we’re here. Once the process gets started, we have a metastable self-perpetuating system of preferences, expectations, and actions capable of persisting indefinitely. As long as uniform conformity is a coordination equilibrium, so that each wants to conform conditionally upon conformity by the others, conforming action produces expectation of conforming action and expectation of conforming action produces conforming action.”

Def 2.6.6. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, in any instance of S among members of P,

Def 2.6.6. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, in any instance of S among members of P,

- everyone conforms to R

Def 2.6.6. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, in any instance of S among members of P,

- everyone conforms to R
- everyone expects everyone else to conform to R

Def 2.6.6. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, in any instance of S among members of P,

- everyone conforms to R
- everyone expects everyone else to conform to R
- everyone prefers to conform to R on condition that the others do, since S is a coordination problem and uniform conformity to R is a proper coordination equilibrium in S."

Def 2.6.6. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, in any instance of S among members of P,

- everyone conforms to R
- everyone expects everyone else to conform to R
- everyone prefers to conform to R on condition that the others do, since S is a coordination problem and uniform conformity to R is a proper coordination equilibrium in S."

A convention consists of a regularity in behavior, a system of mutual expectations, and a system of preferences.

Examples:

Examples:

- Meeting at same place each year.

Examples:

- Meeting at same place each year.
- Oberlin phone convention: original caller calls back.

Examples:

- Meeting at same place each year.
- Oberlin phone convention: original caller calls back.
- Rowing at same speed.

Examples:

- Meeting at same place each year.
- Oberlin phone convention: original caller calls back.
- Rowing at same speed.
- Driving in the right lane.

Examples:

- Meeting at same place each year.
- Oberlin phone convention: original caller calls back.
- Rowing at same speed.
- Driving in the right lane.
- Hunters always doing the same thing.

Examples:

- Meeting at same place each year.
- Oberlin phone convention: original caller calls back.
- Rowing at same speed.
- Driving in the right lane.
- Hunters always doing the same thing.
- Speaking English.

II. Convention Refined

- Common Knowledge

II. Convention Refined

- Common Knowledge
- Alternatives to Conventions

II. Convention Refined

- Common Knowledge
- Alternatives to Conventions
- Degrees of Convention

How does the system of concordant higher-order mutual expectations that fosters coordination arise?

How does the system of concordant higher-order mutual expectations that fosters coordination arise?

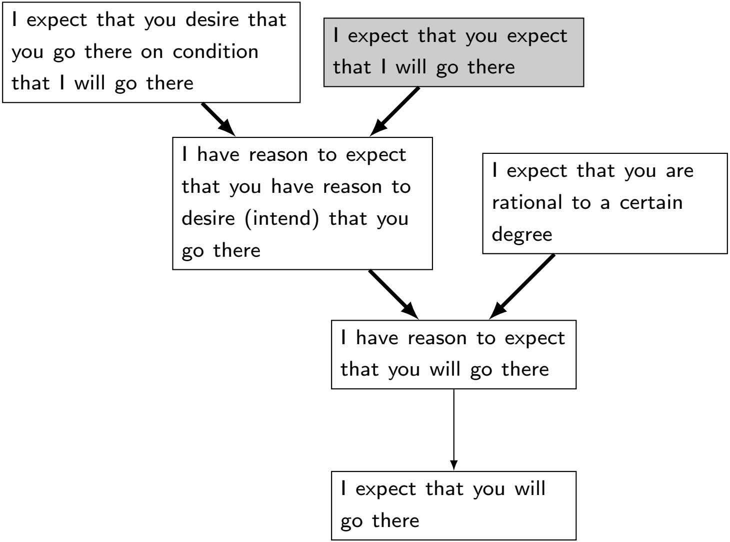

Suppose that we explicitly agree to coordinate. In particular, suppose that the following state of affairs A holds: we have just concluded our meeting in a cafe, we must meet again, and I tell you that I will return to the same cafe tomorrow.

Suppose that we explicitly agree to coordinate. In particular, suppose that the following state of affairs A holds: we have just concluded our meeting in a cafe, we must meet again, and I tell you that I will return to the same cafe tomorrow.

A meets the following conditions:

Suppose that we explicitly agree to coordinate. In particular, suppose that the following state of affairs A holds: we have just concluded our meeting in a cafe, we must meet again, and I tell you that I will return to the same cafe tomorrow.

A meets the following conditions:

(i) You and I have reason to believe that A holds

Suppose that we explicitly agree to coordinate. In particular, suppose that the following state of affairs A holds: we have just concluded our meeting in a cafe, we must meet again, and I tell you that I will return to the same cafe tomorrow.

A meets the following conditions:

(i) You and I have reason to believe that A holds

(ii) A indicates to both of us that you and I have reason to believe that A holds

Suppose that we explicitly agree to coordinate. In particular, suppose that the following state of affairs A holds: we have just concluded our meeting in a cafe, we must meet again, and I tell you that I will return to the same cafe tomorrow.

A meets the following conditions:

(i) You and I have reason to believe that A holds

(ii) A indicates to both of us that you and I have reason to believe that A holds

(iii) A indicates to both of us that I will return

Suppose that we explicitly agree to coordinate. In particular, suppose that the following state of affairs A holds: we have just concluded our meeting in a cafe, we must meet again, and I tell you that I will return to the same cafe tomorrow.

A meets the following conditions:

(i) You and I have reason to believe that A holds

(ii) A indicates to both of us that you and I have reason to believe that A holds

(iii) A indicates to both of us that I will return

where A *indicates* to S that such and such just in case if S had reason to believe that A held then S would thereby have reason to believe that such and such.

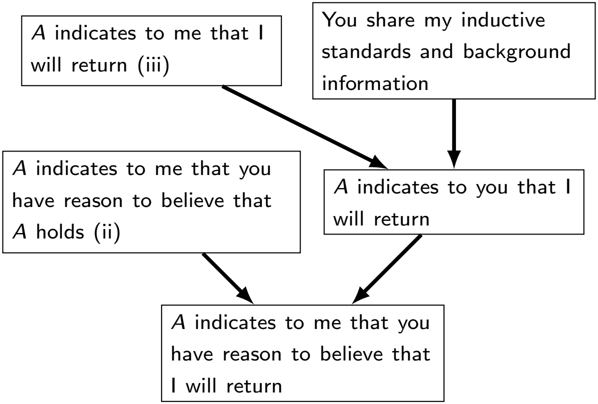

A indicates to me that I will return (iii)
You share my inductive standards and background information
A indicates to me that you have reason to believe that *A* holds (ii)
A indicates to you that I will return
A indicates to me that you have reason to believe that I will return

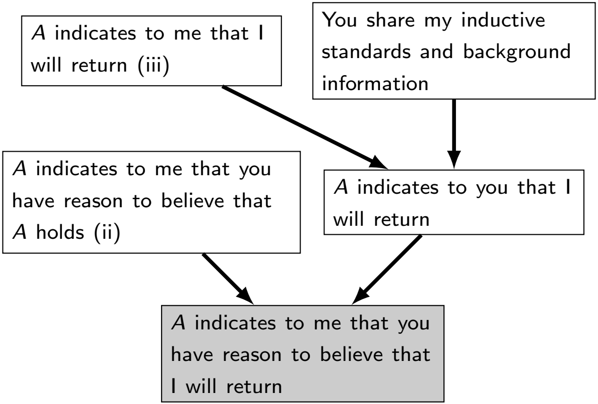
A indicates to me that I will return (iii)
You share my inductive standards and background information
A indicates to me that you have reason to believe that *A* holds (ii)
A indicates to you that I will return
A indicates to me that you have reason to believe that I will return

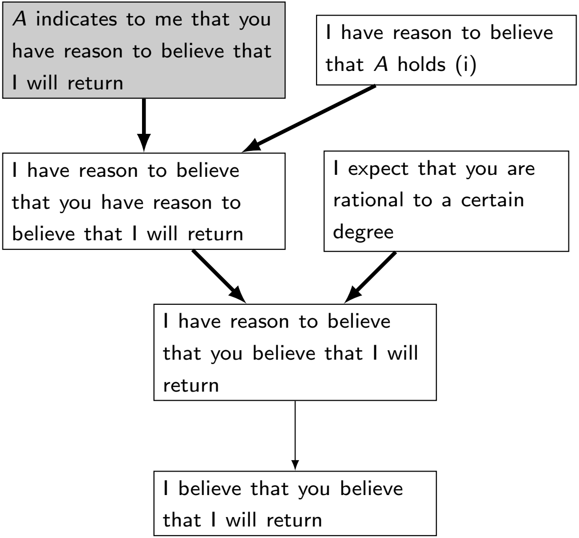

A indicates to me that you have reason to believe that I will return
I have reason to believe that *A* holds (i)
I have reason to believe that you have reason to believe that I will return
I expect that you are rational to a certain degree
I have reason to believe that you believe that I will return
I believe that you believe that I will return

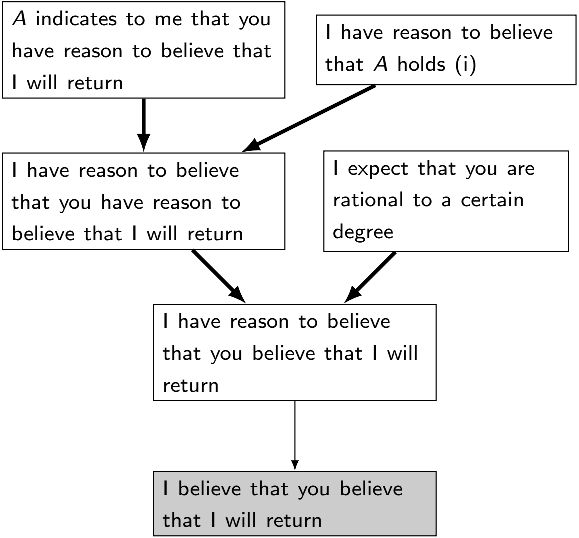
A indicates to me that you have reason to believe that I will return
I have reason to believe that A holds (i)
I have reason to believe that you have reason to believe that I will return
I expect that you are rational to a certain degree
I have reason to believe that you believe that I will return
I believe that you believe that I will return

"I take this example to be typical; all the higher-order expectations involved in sustaining conventions, and more or less all we ever have, seem to be produced in this way."

"I take this example to be typical; all the higher-order expectations involved in sustaining conventions, and more or less all we ever have, seem to be produced in this way."

Def 2.6.7. "It is *common knowledge* in a population P that [such and such] if and only if some state of affairs A holds such that:

- everyone in P has reason to believe that A holds

"I take this example to be typical; all the higher-order expectations involved in sustaining conventions, and more or less all we ever have, seem to be produced in this way."

Def 2.6.7. "It is *common knowledge* in a population P that [such and such] if and only if some state of affairs A holds such that:

- everyone in P has reason to believe that A holds
- A indicates to everyone in P that everyone in P has reason to believe that A holds

"I take this example to be typical; all the higher-order expectations involved in sustaining conventions, and more or less all we ever have, seem to be produced in this way."

Def 2.6.7. "It is *common knowledge* in a population P that [such and such] if and only if some state of affairs A holds such that:

- everyone in P has reason to believe that A holds
- A indicates to everyone in P that everyone in P has reason to believe that A holds
- A indicates to everyone in P that [such and such]."

"I take this example to be typical; all the higher-order expectations involved in sustaining conventions, and more or less all we ever have, seem to be produced in this way."

Def 2.6.7. "It is *common knowledge* in a population P that [such and such] if and only if some state of affairs A holds such that:

- everyone in P has reason to believe that A holds
- A indicates to everyone in P that everyone in P has reason to believe that A holds
- A indicates to everyone in P that [such and such]."

A is a *basis* for common knowledge in P that such and such. Along with mutual ascriptions of rationality, common inductive standards, and background information, A engenders concordant higher-order mutual expectations that such and such is the case.

Past conformity to a convention is a basis for common knowledge about future conformity.

Past conformity to a convention is a basis for common knowledge about future conformity.

Consider a conventional regularity R in population P and let A be the state of affairs that members of P have conformed to R in the past.

Past conformity to a convention is a basis for common knowledge about future conformity.

Consider a conventional regularity R in population P and let A be the state of affairs that members of P have conformed to R in the past.

- everyone in P has reason to believe that A holds

Past conformity to a convention is a basis for common knowledge about future conformity.

Consider a conventional regularity R in population P and let A be the state of affairs that members of P have conformed to R in the past.

- everyone in P has reason to believe that A holds
- A indicates to everyone in P that everyone in P has reason to believe that A holds

Past conformity to a convention is a basis for common knowledge about future conformity.

Consider a conventional regularity R in population P and let A be the state of affairs that members of P have conformed to R in the past.

- everyone in P has reason to believe that A holds
- A indicates to everyone in P that everyone in P has reason to believe that A holds
- A indicates to everyone in P that members of P will conform to R going forward

Def 2.6.8. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, *it is true that, and it is common knowledge in* P *that*, in any instance of S among members of P,

- everyone conforms to R
- everyone expects everyone else to conform to R
- everyone prefers to conform to R on condition that the others do, since S is a coordination problem and uniform conformity to R is a proper coordination equilibrium in S." (emphasis added)

Def 2.6.8. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, *it is true that, and it is common knowledge in* P *that*, in any instance of S among members of P,

- everyone conforms to R
- everyone expects everyone else to conform to R
- everyone prefers to conform to R on condition that the others do, since S is a coordination problem and uniform conformity to R is a proper coordination equilibrium in S." (emphasis added)

There is some state of affairs A that, *inter alia*, indicates that these three conditions hold.

Def 2.6.8. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, *it is true that, and it is common knowledge in* P *that*, in any instance of S among members of P,

- everyone conforms to R
- everyone expects everyone else to conform to R
- everyone prefers to conform to R on condition that the others do, since S is a coordination problem and uniform conformity to R is a proper coordination equilibrium in S." (emphasis added)

There is some state of affairs A that, *inter alia*, indicates that these three conditions hold.

Common knowledge is an important feature of the conventions already discussed.

Def 2.6.8. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, *it is true that, and it is common knowledge in* P *that*, in any instance of S among members of P,

- everyone conforms to R
- everyone expects everyone else to conform to R
- everyone prefers to conform to R on condition that the others do, since S is a coordination problem and uniform conformity to R is a proper coordination equilibrium in S." (emphasis added)

There is some state of affairs A that, *inter alia*, indicates that these three conditions hold.

Common knowledge is an important feature of the conventions already discussed.

Given the common knowledge requirement, certain regularities that do not seem to be conventions will not count as conventions. Example: The condescending driver case.

Common knowledge is not the only source of higher-order expectations.

Common knowledge is not the only source of higher-order expectations.

"Suppose I am a resident of Ableton and I believe everything printed in the *Ableton Argus*. Today's *Argus* prints this story:

The *Bakerville Bugle* is totally unreliable; what it prints is as likely to be false as to be true. Yet the residents of Bakerville believe everything in it. Today's *Bugle* printed this story:

The *Charlie City Crier* is totally unreliable; what it prints is as likely to be false as to be true. Yet the residents of Charlie City believe everything in it. Today's *Crier* printed this story:

The *Dogpatch Daily* is totally unreliable; what it prints is as likely to be false as to be true. Yet the residents of Dogpatch believe everything in it. Today's *Daily* printed this story:

Tomorrow it will rain cats and dogs."

Common knowledge is not the only source of higher-order expectations.

"Suppose I am a resident of Ableton and I believe everything printed in the *Ableton Argus*. Today's *Argus* prints this story:

The *Bakerville Bugle* is totally unreliable; what it prints is as likely to be false as to be true. Yet the residents of Bakerville believe everything in it. Today's *Bugle* printed this story:

The *Charlie City Crier* is totally unreliable; what it prints is as likely to be false as to be true. Yet the residents of Charlie City believe everything in it. Today's *Crier* printed this story:

The *Dogpatch Daily* is totally unreliable; what it prints is as likely to be false as to be true. Yet the residents of Dogpatch believe everything in it. Today's *Daily* printed this story:

Tomorrow it will rain cats and dogs."

I expect the residents of Bakerville to expect the residents of Charlie City to expect the residents of Dogpatch to expect that it will rain cats and dogs. But I do not have lower-order expectations.

The second refinement of the analysis of convention is motivated by regularities in action that cannot be broken down into self-contained coordination problems.

The second refinement of the analysis of convention is motivated by regularities in action that cannot be broken down into self-contained coordination problems.

Ex. Oligopoly.

"Suppose that we are contented oligopolists. As the price of our raw material varies, we must each set new prices. It is to no one's advantage to set his prices higher than the others set theirs, since if he does he tends to lose his share of the market. Nor is it to anyone's advantage to set his prices lower than the others set theirs, since if he does he menaces his competitors and incurs their retaliation. So each must set his prices within the range of prices he expects the others to set."

The second refinement of the analysis of convention is motivated by regularities in action that cannot be broken down into self-contained coordination problems.

Ex. Oligopoly.

"Suppose that we are contented oligopolists. As the price of our raw material varies, we must each set new prices. It is to no one's advantage to set his prices higher than the others set theirs, since if he does he tends to lose his share of the market. Nor is it to anyone's advantage to set his prices lower than the others set theirs, since if he does he menaces his competitors and incurs their retaliation. So each must set his prices within the range of prices he expects the others to set."

Now, consider a convention where we follow a price leader—that is, one of us initiates price changes in a way that suits all of us. There is a regularity in action. But to think of us as repeatedly facing self-contained coordination problems is precious. Each of us can change our prices whenever we want.

Def 2.6.9. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, it is true that, and it is common knowledge in P that, in any instance of S among members of P,

Def 2.6.9. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, it is true that, and it is common knowledge in P that, in any instance of S among members of P,

- everyone conforms to R

Def 2.6.9. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, it is true that, and it is common knowledge in P that, in any instance of S among members of P,

- everyone conforms to R
- everyone expects everyone else to conform to R

Def 2.6.9. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, it is true that, and it is common knowledge in P that, in any instance of S among members of P,

- everyone conforms to R
- everyone expects everyone else to conform to R
- everyone has approximately the same preferences regarding all possible combinations of actions

Def 2.6.9. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, it is true that, and it is common knowledge in P that, in any instance of S among members of P,

- everyone conforms to R
- everyone expects everyone else to conform to R
- everyone has approximately the same preferences regarding all possible combinations of actions
- everyone prefers that everyone conform to R, on condition that at least all but one conform to R

Def 2.6.9. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, it is true that, and it is common knowledge in P that, in any instance of S among members of P,

- everyone conforms to R
- everyone expects everyone else to conform to R
- everyone has approximately the same preferences regarding all possible combinations of actions
- everyone prefers that everyone conform to R, on condition that at least all but one conform to R
- everyone would prefer that everyone conform to R', on condition that at least all but one conform to R'

where R' is some possible regularity in the behavior of members of P in S, such that no one in any instance of S among members of P could conform both to R' and to R."

- everyone has approximately the same preferences regarding all possible combinations of actions
- everyone prefers that everyone conform to R, on condition that at least all but one conform to R
- everyone would prefer that everyone conform to R', on condition that at least all but one conform to R'

- everyone has approximately the same preferences regarding all possible combinations of actions
- everyone prefers that everyone conform to R, on condition that at least all but one conform to R
- everyone would prefer that everyone conform to R', on condition that at least all but one conform to R'

If S is a self-contained interactive choice situation, then the first condition ensures that it is a game of coordination rather than a game of conflict.

- everyone has approximately the same preferences regarding all possible combinations of actions
- everyone prefers that everyone conform to R, on condition that at least all but one conform to R
- everyone would prefer that everyone conform to R', on condition that at least all but one conform to R'

If S is a self-contained interactive choice situation, then the first condition ensures that it is a game of coordination rather than a game of conflict.

If S is a self-contained interactive choice situation, then the second condition ensures that uniform conformity to R is a proper coordination equilibrium.

- everyone has approximately the same preferences regarding all possible combinations of actions
- everyone prefers that everyone conform to R, on condition that at least all but one conform to R
- everyone would prefer that everyone conform to R', on condition that at least all but one conform to R'

If S is a self-contained interactive choice situation, then the first condition ensures that it is a game of coordination rather than a game of conflict.

If S is a self-contained interactive choice situation, then the second condition ensures that uniform conformity to R is a proper coordination equilibrium.

If S is a self-contained interactive choice situation, then the third condition ensures that uniform conformity to R' is another proper coordination equilibrium.

- everyone has approximately the same preferences regarding all possible combinations of actions
- everyone prefers that everyone conform to R, on condition that at least all but one conform to R
- everyone would prefer that everyone conform to R', on condition that at least all but one conform to R'

If S is a self-contained interactive choice situation, then the first condition ensures that it is a game of coordination rather than a game of conflict.

If S is a self-contained interactive choice situation, then the second condition ensures that uniform conformity to R is a proper coordination equilibrium.

If S is a self-contained interactive choice situation, then the third condition ensures that uniform conformity to R' is another proper coordination equilibrium.

Thus, the new definition generalizes the old definition by removing the game-theoretic scaffolding.

If R is a convention, then R' might have been our convention instead. Conventions are inherently *arbitrary*.

If R is a convention, then R' might have been our convention instead. Conventions are inherently *arbitrary*.

Since convention requires common knowledge of this arbitrariness, whether a regularity counts as a convention can be sensitive to exposure and open-mindedness to alternatives:

"What is not conventional among narrow-minded and inflexible people, who would not know what to do if others began to behave differently, may be conventional among more adaptable people. What is not conventional may become conventional when news arrives of aliens who behave differently; or when somebody invents a new way of behaving, even a new way no one adopts. When children and the feeble-minded conform to our conventions, they may not take part in them *as* conventions, for they may lack any conditional preference for conformity to an alternative; or they may have the proper preferences, but not as an item of common knowledge. I find these corollaries of our analysis of convention neither welcome nor unwelcome. The analysis is settling questions hitherto left open."

The third and final refinement of the analysis of convention makes things less strict.

The third and final refinement of the analysis of convention makes things less strict.

Def 2.6.10. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, it is true that, and it is common knowledge in P that, in *almost* any instance of S among members of P,

The third and final refinement of the analysis of convention makes things less strict.

Def 2.6.10. “A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, it is true that, and it is common knowledge in P that, in *almost* any instance of S among members of P,

- *almost* everyone conforms to R

The third and final refinement of the analysis of convention makes things less strict.

Def 2.6.10. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, it is true that, and it is common knowledge in P that, in *almost* any instance of S among members of P,

- *almost* everyone conforms to R
- *almost* everyone expects *almost* everyone else to conform to R

The third and final refinement of the analysis of convention makes things less strict.

Def 2.6.10. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, it is true that, and it is common knowledge in P that, in *almost* any instance of S among members of P,

- *almost* everyone conforms to R
- *almost* everyone expects *almost* everyone else to conform to R
- *almost* everyone has approximately the same preferences regarding all possible combinations of actions

The third and final refinement of the analysis of convention makes things less strict.

Def 2.6.10. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, it is true that, and it is common knowledge in P that, in *almost* any instance of S among members of P,

- *almost* everyone conforms to R
- *almost* everyone expects *almost* everyone else to conform to R
- *almost* everyone has approximately the same preferences regarding all possible combinations of actions
- *almost* everyone prefers that *any one more* conform to R, on condition that *almost everyone* conform to R

The third and final refinement of the analysis of convention makes things less strict.

Def 2.6.10. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, it is true that, and it is common knowledge in P that, in *almost* any instance of S among members of P,

- *almost* everyone conforms to R
- *almost* everyone expects *almost* everyone else to conform to R
- *almost* everyone has approximately the same preferences regarding all possible combinations of actions
- *almost* everyone prefers that *any one more* conform to R, on condition that *almost everyone* conform to R
- *almost* everyone would prefer that *any one more* conform to R', on condition that *almost everyone* conform to R'

where R' is some possible regularity in the behavior of members of P in S, such that *almost* no one in *almost* any instance of S among members of P could conform both to R' and to R." (my emphasis)

IV. Convention and Communication

- Sample Signals

IV. Convention and Communication

- Sample Signals
- Analysis of Signaling

IV. Convention and Communication

- Sample Signals
- Analysis of Signaling
- Meaning of Signals

Ex. Paul Revere.

The sexton of the Old North Church wants to communicate information about the British army to Paul Revere.

Both the sexton and Paul Revere must choose a contingency plan that will guide their behavior going forward. Each will choose his plan with regard to his expectation of the other's choice.

Some candidate plans for the sexton:

Some candidate plans for the sexton:

Plan R1:
If the redcoats are observed staying home, hang no lantern.
If the redcoats are observed setting out by land, hang one lantern.
If the redcoats are observed setting out by sea, hang two lanterns.

Some candidate plans for the sexton:

Plan R1:
If the redcoats are observed staying home, hang no lantern.
If the redcoats are observed setting out by land, hang one lantern.
If the redcoats are observed setting out by sea, hang two lanterns.

Plan R2:
If the redcoats are observed staying home, hang one lantern.
If the redcoats are observed setting out by land, hang two lanterns.
If the redcoats are observed setting out by sea, hang no lantern.

Some candidate plans for the sexton:

Plan R1:
If the redcoats are observed staying home, hang no lantern.
If the redcoats are observed setting out by land, hang one lantern.
If the redcoats are observed setting out by sea, hang two lanterns.

Plan R2:
If the redcoats are observed staying home, hang one lantern.
If the redcoats are observed setting out by land, hang two lanterns.
If the redcoats are observed setting out by sea, hang no lantern.

Plan R3:
If the redcoats are observed staying home, hang one lantern.
If the redcoats are observed setting out by land, hang no lantern.
If the redcoats are observed setting out by sea, hang two lanterns.

Some candidate plans for Revere:

Some candidate plans for Revere:

Plan C1:
If no lantern is observed in the belfry, go home.
If one lantern is observed in the belfry, warn the countryside that the redcoats are coming by land.
If two lanterns are observed in the belfry, warn the countryside that the redcoats are coming by sea.

Some candidate plans for Revere:

Plan C1:
If no lantern is observed in the belfry, go home.
If one lantern is observed in the belfry, warn the countryside that the redcoats are coming by land.
If two lanterns are observed in the belfry, warn the countryside that the redcoats are coming by sea.

Plan C2:
If no lantern is observed in the belfry, warn the countryside that the redcoats are coming by sea.
If one lantern is observed in the belfry, go home.
If two lanterns are observed in the belfry, warn the countryside that the redcoats are coming by land.

Some candidate plans for Revere:

Plan C1:
If no lantern is observed in the belfry, go home.
If one lantern is observed in the belfry, warn the countryside that the redcoats are coming by land.
If two lanterns are observed in the belfry, warn the countryside that the redcoats are coming by sea.

Plan C2:
If no lantern is observed in the belfry, warn the countryside that the redcoats are coming by sea.
If one lantern is observed in the belfry, go home.
If two lanterns are observed in the belfry, warn the countryside that the redcoats are coming by land.

Plan C3:
If no lantern is observed in the belfry, warn the countryside that the redcoats are coming by land.
If one lantern is observed in the belfry, go home.
If two lanterns are observed in the belfry, warn the countryside that the redcoats are coming by sea.

The choice of contingency plans is a coordination problem.

	C1	C2	C3
R1	2,2	0,0	1,1
R2	0,0	2,2	1,1
R3	1,1	1,1	2,2

The choice of contingency plans is a coordination problem.

	C1	C2	C3
R1	2,2	0,0	1,1
R2	0,0	2,2	1,1
R3	1,1	1,1	2,2

In reality, the proper coordination equilibrium $\langle R1, C1 \rangle$ was reached through explicit agreement. The sexton and Paul Revere agreed upon signals for a single occasion.

“I have now described the character of a case of signaling without mentioning the meaning of the signals: that two lanterns meant that the redcoats were coming by sea, or whatever. But nothing important seems to have been left unsaid, so what has been said must somehow imply that the signals have their meanings.”

Examples of signaling conventions:

Examples of signaling conventions:

- International Code of Signals (for ships).

Examples of signaling conventions:

- International Code of Signals (for ships).
- Helping a truck park.

Examples of signaling conventions:

- International Code of Signals (for ships).
- Helping a truck park.
- Blazing a trail.

Def 2.6.11. A (two-sided) *signaling problem* is a situation S involving a *communicator* and *audience* such that it is common knowledge among them that:

Def 2.6.11. A (two-sided) *signaling problem* is a situation S involving a *communicator* and *audience* such that it is common knowledge among them that:

- Exactly one of the *states* $s_1, ..., s_m$ holds and the communicator is well-positioned to know which state holds

Def 2.6.11. A (two-sided) *signaling problem* is a situation S involving a *communicator* and *audience* such that it is common knowledge among them that:

- Exactly one of the *states* $s_1, ..., s_m$ holds and the communicator is well-positioned to know which state holds
- Each member of the audience can perform one of the *responses* $r_1, ..., r_m$

Def 2.6.11. A (two-sided) *signaling problem* is a situation S involving a *communicator* and *audience* such that it is common knowledge among them that:

- Exactly one of the *states* $s_1, ..., s_m$ holds and the communicator is well-positioned to know which state holds
- Each member of the audience can perform one of the *responses* $r_1, ..., r_m$
- There is a 1-1 function $F : \{s_i\} \mapsto \{r_i\}$ from states to responses such that everyone prefers that each member of the audience do $F(s_i)$ on condition that s_i holds

  (a function $f(x)$ is 1-1 just in case $f(x) \neq f(y)$ whenever $x \neq y$)

Def 2.6.11. A (two-sided) *signaling problem* is a situation S involving a *communicator* and *audience* such that it is common knowledge among them that:

- Exactly one of the *states* $s_1, ..., s_m$ holds and the communicator is well-positioned to know which state holds
- Each member of the audience can perform one of the *responses* $r_1, ..., r_m$
- There is a 1-1 function $F : \{s_i\} \mapsto \{r_i\}$ from states to responses such that everyone prefers that each member of the audience do $F(s_i)$ on condition that s_i holds

  (a function $f(x)$ is 1-1 just in case $f(x) \neq f(y)$ whenever $x \neq y$)
- The communicator can send one of the *signals* $\sigma_1, ..., \sigma_n$ $(n \geq m)$ and the audience is well-positioned to know which signal was sent

Def 2.6.12. A *communicator's contingency plan* $F_c : \{s_i\} \mapsto \{\sigma_i\}$ is a function from states to signals. If F_c is a 1-1 function, then this plan is *admissible*.

Def 2.6.12. A *communicator's contingency plan* $F_c : \{s_i\} \mapsto \{\sigma_i\}$ is a function from states to signals. If F_c is a 1-1 function, then this plan is *admissible*.

Def 2.6.13. An *audience's contingency plan* $F_a : \{\sigma_i\} \mapsto \{r_i\}$ is a 1-1 function from signals to responses. If the ranges of F_a and F coincide, then F_a is *admissible*.

Def 2.6.12. A *communicator's contingency plan* $F_c : \{s_i\} \mapsto \{\sigma_i\}$ is a function from states to signals. If F_c is a 1-1 function, then this plan is *admissible*.

Def 2.6.13. An *audience's contingency plan* $F_a : \{\sigma_i\} \mapsto \{r_i\}$ is a 1-1 function from signals to responses. If the ranges of F_a and F coincide, then F_a is *admissible*.

The composition $f \circ g$ is the function $f(g(x))$.

Def 2.6.12. A *communicator's contingency plan* $F_c : \{s_i\} \mapsto \{\sigma_i\}$ is a function from states to signals. If F_c is a 1-1 function, then this plan is *admissible*.

Def 2.6.13. An *audience's contingency plan* $F_a : \{\sigma_i\} \mapsto \{r_i\}$ is a 1-1 function from signals to responses. If the ranges of F_a and F coincide, then F_a is *admissible*.

The composition $f \circ g$ is the function $f(g(x))$.

Def 2.6.14. A *signaling system* $\langle F_c, F_a \rangle$ is a system of communicator's and audience's contingency plans where $F_a \circ F_c = F$.

s_1	r_1
s_2	r_2
s_3	r_3
s_4	r_4
s_5	r_5
$\vdots$	$\vdots$
s_m	r_m
states	**responses**

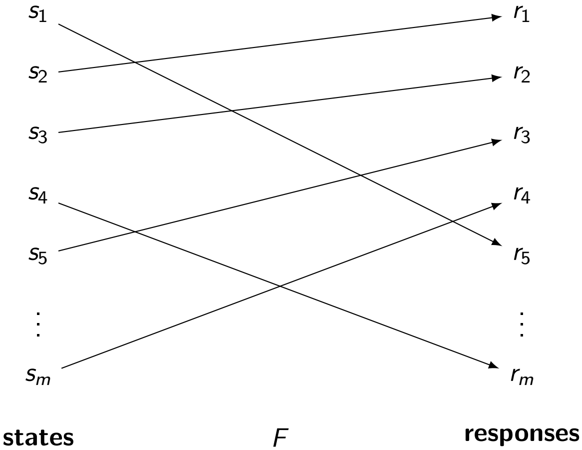
s_1
s_2
s_3
s_4
s_5
s_m
r_1
r_2
r_3
r_4
r_5
r_m
states
F
responses

s_1	σ_1	r_1
s_2	σ_2	r_2
s_3	σ_3	r_3
s_4	σ_4	r_4
s_5	σ_5	r_5
$\vdots$	$\vdots$	$\vdots$
s_m	σ_n	r_m
states	**signals**	**responses**

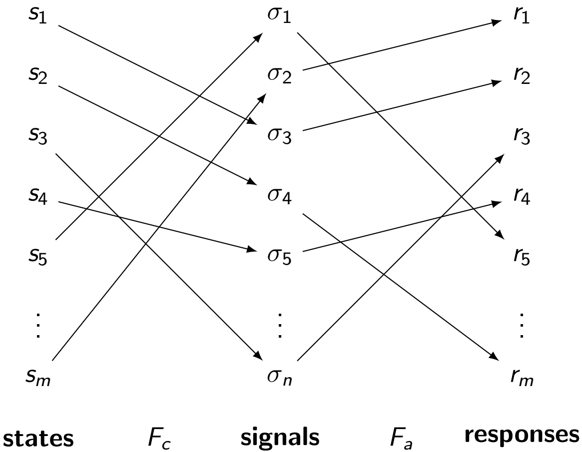

s_1
s_2
s_3
s_4
s_5
s_m
σ_1
σ_2
σ_3
σ_4
σ_5
σ_n
r_1
r_2
r_3
r_4
r_5
r_m
states
F_c
signals
F_a
responses

Ex. Paul Revere.

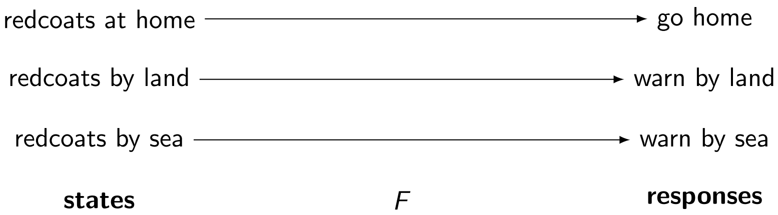

Ex. Paul Revere.

states	R1	signals	C1	responses
redcoats at home	⟶	hang no lantern	⟶	go home
redcoats by land	⟶	hang one lantern	⟶	warn by land
redcoats by sea	⟶	hang two lanterns	⟶	warn by sea

Ex. Paul Revere.

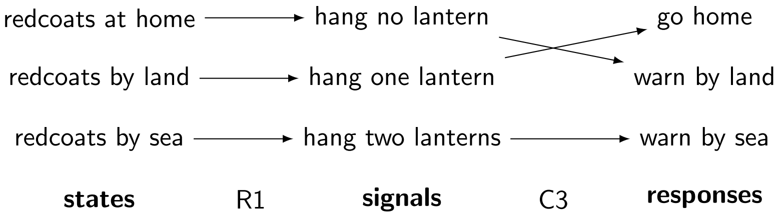

Ex. Paul Revere.

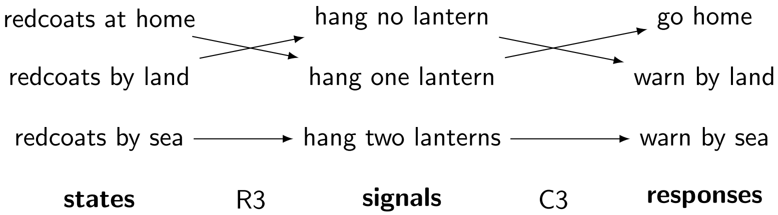

Fact. All and only admissible contingency plans belong to signaling systems.

Fact. All and only admissible contingency plans belong to signaling systems.

Fact. In a signaling problem with m states and n signals, there are $\frac{n!}{(n-m)!}$ signaling systems.

Fact. All and only admissible contingency plans belong to signaling systems.

Fact. In a signaling problem with m states and n signals, there are $\frac{n!}{(n-m)!}$ signaling systems.

In Paul Revere, there are $\frac{3!}{0!} = 6$ signaling systems.

Fact. All and only admissible contingency plans belong to signaling systems.

Fact. In a signaling problem with m states and n signals, there are $\frac{n!}{(n-m)!}$ signaling systems.

In Paul Revere, there are $\frac{3!}{0!} = 6$ signaling systems.

The choice between contingency plans is a coordination problem and signaling systems are proper coordination equilibria (there may be other coordination equilibria besides).

Fact. All and only admissible contingency plans belong to signaling systems.

Fact. In a signaling problem with m states and n signals, there are $\frac{n!}{(n-m)!}$ signaling systems.

In Paul Revere, there are $\frac{3!}{0!} = 6$ signaling systems.

The choice between contingency plans is a coordination problem and signaling systems are proper coordination equilibria (there may be other coordination equilibria besides).

Def 2.6.15. A *signaling convention* is any convention whereby communicators and audiences faced with a signaling problem S do their part of a signaling system $\langle F_c, F_a \rangle$. This system is a *conventional signaling system*.

Def 2.6.16. A *verbal expression* is a finite sequence of types of vocal sounds or marks.

Def 2.6.16. A *verbal expression* is a finite sequence of types of vocal sounds or marks.

Def 2.6.17. A *verbal signal* is an act of uttering or inscribing a verbal expression.

Def 2.6.16. A *verbal expression* is a finite sequence of types of vocal sounds or marks.

Def 2.6.17. A *verbal signal* is an act of uttering or inscribing a verbal expression.

Much of our language falls under verbal signaling.

“If we endow a hypothetical community with a great many verbal signaling conventions for use in various activities, with verbal expressions suitably chosen *ad hoc*, we shall be able to simulate a community of language users—say, ourselves—rather well. An observer who stayed in the background watching these people use conventional verbal signals as they went about their business might take a long time to realize that they were not ordinary language users. But an observer who tried to converse with them would notice some deficiencies. He would find that every verbal expression they used was conventionally associated with some readily observable state of affairs, or with some definite responsive action, or both. And he would find that they could use only finitely many verbal expressions, so that the conventions governing their verbal signaling could be described by mentioning each expression used.”

"If we endow a hypothetical community with a great many verbal signaling conventions for use in various activities, with verbal expressions suitably chosen *ad hoc*, we shall be able to simulate a community of language users—say, ourselves—rather well. An observer who stayed in the background watching these people use conventional verbal signals as they went about their business might take a long time to realize that they were not ordinary language users. But an observer who tried to converse with them would notice some deficiencies. He would find that every verbal expression they used was conventionally associated with some readily observable state of affairs, or with some definite responsive action, or both. And he would find that they could use only finitely many verbal expressions, so that the conventions governing their verbal signaling could be described by mentioning each expression used."

"Yet it remains true that our hypothetical verbal signalers do not do anything we do not do. We just do more. Their use of language duplicates a fragment of ours."

At this point, we can introduce *meaning* into the picture.

At this point, we can introduce *meaning* into the picture.

Consider a signaling system $\langle F_c, F_a \rangle$ where $F_c(s) = \sigma$ and $F_a(\sigma) = r$.

At this point, we can introduce *meaning* into the picture.

Consider a signaling system $\langle F_c, F_a \rangle$ where $F_c(s) = \sigma$ and $F_a(\sigma) = r$.

On the one hand, we might say that σ *means* in $\langle F_c, F_a \rangle$ that s holds.

At this point, we can introduce *meaning* into the picture.

Consider a signaling system $\langle F_c, F_a \rangle$ where $F_c(s) = \sigma$ and $F_a(\sigma) = r$.

On the one hand, we might say that σ *means* in $\langle F_c, F_a \rangle$ that s holds.

On the other hand, we might say that σ *means* in $\langle F_c, F_a \rangle$ to do r.

At this point, we can introduce *meaning* into the picture.

Consider a signaling system $\langle F_c, F_a \rangle$ where $F_c(s) = \sigma$ and $F_a(\sigma) = r$.

On the one hand, we might say that σ *means* in $\langle F_c, F_a \rangle$ that s holds.

On the other hand, we might say that σ *means* in $\langle F_c, F_a \rangle$ to do r.

In the former case, σ is an *indicative* signal.

In the latter case, σ is an *imperative* signal.

A signal is *neutral* if it is equally properly called an indicative and imperative signal.

At this point, we can introduce *meaning* into the picture.

Consider a signaling system $\langle F_c, F_a \rangle$ where $F_c(s) = \sigma$ and $F_a(\sigma) = r$.

On the one hand, we might say that σ *means* in $\langle F_c, F_a \rangle$ that s holds.

On the other hand, we might say that σ *means* in $\langle F_c, F_a \rangle$ to do r.

In the former case, σ is an *indicative* signal.

In the latter case, σ is an *imperative* signal.

A signal is *neutral* if it is equally properly called an indicative and imperative signal.

Whether a signal counts as indicative, imperative, or neutral depends on whether the communicator or audience has to engage in significant deliberation in the signaling system.

We can also introduce *truth* into the picture.

We can also introduce *truth* into the picture.

If σ is an indicative signal that s holds in system $\langle F_c, F_a \rangle$, then we might say that σ is *true in* $\langle F_c, F_a \rangle$ when s holds and σ is *false in* $\langle F_c, F_a \rangle$ when s does not hold.

We can also introduce *truth* into the picture.

If σ is an indicative signal that s holds in system $\langle F_c, F_a \rangle$, then we might say that σ is *true in* $\langle F_c, F_a \rangle$ when s holds and σ is *false in* $\langle F_c, F_a \rangle$ when s does not hold.

Since signals are actions, we are ascribing truth/falsity to actions. But in the case of verbal signals, we might also ascribe truth/falsity to verbal expressions or their tokenings.

We can also introduce *truth* into the picture.

If σ is an indicative signal that s holds in system $\langle F_c, F_a \rangle$, then we might say that σ is *true in* $\langle F_c, F_a \rangle$ when s holds and σ is *false in* $\langle F_c, F_a \rangle$ when s does not hold.

Since signals are actions, we are ascribing truth/falsity to actions. But in the case of verbal signals, we might also ascribe truth/falsity to verbal expressions or their tokenings.

To engage in a conventional signaling system $\langle F_c, F_a \rangle$ is to follow a convention of *truthfulness in* $\langle F_c, F_a \rangle$.

We can also introduce *truth* into the picture.

If σ is an indicative signal that s holds in system $\langle F_c, F_a \rangle$, then we might say that σ is *true in* $\langle F_c, F_a \rangle$ when s holds and σ is *false in* $\langle F_c, F_a \rangle$ when s does not hold.

Since signals are actions, we are ascribing truth/falsity to actions. But in the case of verbal signals, we might also ascribe truth/falsity to verbal expressions or their tokenings.

To engage in a conventional signaling system $\langle F_c, F_a \rangle$ is to follow a convention of *truthfulness in* $\langle F_c, F_a \rangle$.

"In any instance of S among members of P, the communicator tries to give whichever signal is true *under the prevailing convention* in that instance, and the audience responds by doing whatever seems best on the assumption that he has succeeded in so doing...We can say that a signal σ is *true* in P in any instance of S if and only if there is some suitable signaling system that is conventionally adopted in P and σ is true in that signaling system in that instance of S."

If $\langle F_c, F_a \rangle$ is a verbal signaling system, then we can think of this system as a *language* $\mathcal{L}$.

If $\langle F_c, F_a \rangle$ is a verbal signaling system, then we can think of this system as a *language* $\mathcal{L}$.

Speakers of this language follow a convention of *truthfulness in* $\mathcal{L}$.