

A **model of decisions under uncertainty** is characterized by:
a set of alternative choices C , a set of possible states of the world S ,
a utility function $u:C \times S \rightarrow \mathbb{R}$, and a probability distribution p in $\Delta(S)$.
Suppose that C and S are nonempty finite sets.

Here we use the notation $\Delta(S) = \{q \in \mathbb{R}^S \mid q(s) \geq 0 \forall s, \sum_{\theta \in S} q(\theta) = 1\}$.

The expected utility hypothesis says that an optimal decision should maximize expected utility $Eu(c) = Eu(c|p) = \sum_{\theta \in S} p(\theta)u(c,\theta)$ over all c in C , for some utility function u that is appropriate for the decision maker.

Example 1. Consider an example with choices $C = \{T,M,B\}$, state $S = \{L,R\}$, and

$u(c,s)$:	L	R
T	7	2
M	2	7
B	5	6

To describe the probability distribution parametrically, let r be the probability of state R.

So $Eu(T) = 7(1-r)+2r$, $Eu(M) = 2(1-r)+7r$, $Eu(B) = 5(1-r)+6r$.

Then B is optimal when $5(1-r)+6r \geq 7(1-r)+2r$ and $5(1-r)+6r \geq 2(1-r)+7r$, which are satisfied when $1/3 \leq r \leq 3/4$.

T is optimal when $r \leq 1/3$. M is optimal when $r \geq 3/4$.

Fact: Given the utility function $u:C \times S \rightarrow \mathbb{R}$ and some choice option $d \in C$, the set of probability distributions that make d optimal is a closed convex (possibly empty) subset of $\Delta(S)$.

This set (of probabilities in $\Delta(S)$ that make d optimal) is empty if and only if there exists some randomized strategy σ in $\Delta(C)$ such that $u(d,s) < \sum_{c \in C} \sigma(c)u(c,s) \forall s \in S$.

When these inequalities hold, we say that d is strongly dominated by σ .

[Proof: $\{x \in \mathbb{R}^S \mid \exists \sigma \in \Delta(C) \text{ s.t. } x_s \leq \sum_{c \in C} \sigma(c)u(c,s) \forall s\}$ is a convex subset of \mathbb{R}^S . d is strongly dominated iff $(u(d,s))_{s \in S}$ is in its interior. Use supporting-hyperplane thm, MWG p. 949.]

Example 2: As above, $C = \{T,M,B\}$, $S = \{L,R\}$, and u is same except $u(B,R) = 3$.

$u(c,s)$:	L	R
T	7	2
M	2	7
B	5	3

As before, B would be the second-best choice in either state (if the state were known).

B would be an optimal decision under uncertainty when

$5(1-r)+3r \geq 7(1-r)+2r$ and $5(1-r)+3r \geq 2(1-r)+7r$,

which are satisfied when $r \geq 2/3$ and $3/7 \geq r$, which is impossible! So B cannot be optimal.

T is optimal when $r \leq 1/2$. M is optimal when $r \geq 1/2$.

Now consider a randomized strategy that chooses T with some probability $\sigma(T)$

and chooses M otherwise, with probability $\sigma(M) = 1 - \sigma(T)$.

B would be strongly dominated by this randomized strategy σ if

$5 < \sigma(T)7 + (1 - \sigma(T))2$ (B worse than σ in state L), and

$3 < \sigma(T)2 + (1 - \sigma(T))7$ (B worse than σ in state R).

These inequalities are satisfied when $3/5 < \sigma(T) < 4/5$. For example, $\sigma(T) = 0.7$ works. That is, B is strongly dominated by $0.7[T] + 0.3[M]$, as $5 < 0.7 \times 7 + 0.3 \times 2 = 5.5$ and $3 < 0.7 \times 2 + 0.3 \times 7 = 3.5$.

Computing randomized Nash equilibria for games that are larger than 2×2 can be difficult, but working a few examples can help you better understand Nash's subtle concept of equilibrium. At the top of page 138, Osborne describes a general procedure for finding randomized Nash equilibria for any finite game, based on the characterization in Proposition 116.2. Here, we describe this procedure in somewhat different terms, with an illustrative application.

We are given some game, including a given set of players N and, for each i in N , a given set of feasible actions A_i for player i and a given payoff function $u_i: A_1 \times \dots \times A_n \rightarrow \mathbb{R}$ for player i .

The support of a randomized equilibrium is, for each player, the set of actions that have positive probability of being chosen in this equilibrium.

To find a Nash equilibrium, we can apply the following 4-step method:

(1) Guess a support for all players. That is, for each player i , let S_i be a subset of i 's actions A_i , and let us guess that S_i is the set of actions that player i will use with positive probability.

(2) Consider the smaller game where the action set for each player i is reduced to S_i , and try to find an equilibrium where all of these actions get positive probability.

To do this, we need to solve a system of equations for some unknown quantities.

The unknowns: For each player i in N and each action s_i in i 's support S_i , let $\sigma_i(s_i)$ denote i 's probability of choosing s_i , and let w_i denote player i 's expected payoff in the equilibrium.

The equations: For each player i , the sum of these probabilities $\sigma_i(s_i)$ must equal 1.

For each player i and each action s_i in S_i , player i 's expected payoff when he chooses s_i but all other players randomize independently according to their σ_j probabilities must be equal to w_i .

Let $Eu_i(a_i | \sigma_{-i})$ denote player i 's expected payoff when he chooses action a_i and all other players are expected to randomize independently according to their σ_j probabilities.

Then the equations can be written: $\sum_{s_i \in S_i} \sigma_i(s_i) = 1 \quad \forall i \in N$; and $Eu_i(s_i | \sigma_{-i}) = w_i \quad \forall i \in N \quad \forall s_i \in S_i$.

(Here \forall means "for all", \in means "in".) We have as many equations as unknowns ($w_i, \sigma_i(s_i)$).

(For any action a_i that is **not** in player i 's support S_i , we have $\sigma_i(a_i) = 0$.)

If the equations in step 2 have no solution, then we guessed the wrong support, and so we must return to step 1 and guess a new support.

Assuming that we have a solution from step (2), continue to (3) and (4)

(3) The solution from (2) would be nonsense if any of the "probabilities" were negative.

That is, for every player i in N and every action s_i in i 's support S_i , we need $\sigma_i(s_i) \geq 0$.

If these nonnegativity conditions are not satisfied by a solution, then we have not found an equilibrium with the guessed support, and so we must return to step 1 and guess a new support.

If we have a solution that satisfies all these nonnegativity conditions, then it is a randomized equilibrium of the reduced game where each player must can only choose actions in S_i .

(4) A solution from (2) that satisfies the condition in (3) would still not be an equilibrium of the original game, however, if any player would prefer an action outside the guessed support.

So next we must ask, for each player i and for each action a_i that is in A_i but is not in the guessed support S_i , could player i do better than w_i by choosing a_i when all other players randomize independently according to their σ_j probabilities? Recall $Eu_i(s_i | \sigma_{-i}) = w_i$ for all s_i in S_i .

So now, for every action a_i that is in A_i but is not in S_i , we need $Eu_i(a_i | \sigma_{-i}) \leq w_i$.

If our solution satisfies all these inequalities then it is an equilibrium of the given game.

But if any of these inequalities is violated (some $Eu_i(a_i | \sigma_{-i}) > w_i$), then we have not found an equilibrium with the guessed support, and so we must return to step 1 and guess a new support.

In a finite game, there are only a finite number of possible supports to consider.

Thus in equilibrium, w_i ($\forall i \in N$) and $\sigma_i(a_i)$ ($\forall i \in N, \forall a_i \in A_i$) must satisfy: $\sum_{a_i} \sigma_i(a_i) = 1 \quad \forall i \in N$, and $\sigma_i(a_i) \geq 0$ and $Eu_i(a_i | \sigma_{-i}) \leq w_i$ with at least one equality $\forall i \in N \quad \forall a_i \in A_i$ (*complementary slackness*).

Example. Find all Nash equilibria (pure and mixed) of the following 2×3 game:

		Player 2		
		L	M	R
Player 1	T	7, 2	2, 7	3, 6
	B	2, 7	7, 2	4, 5

It is easy to see that this game has no pure-strategy equilibria (2's best response to T is M, but T is not 1's best response to M; and 2's best response to B is L, but B is not 1's best response to L).

This eliminates the six cases where each player's support is just one action.

Furthermore, when either player is restricted to just one action, the other player always has a unique best response, and so there are no equilibria where only one player randomizes.

That is, both players must have at least two actions in the support of any equilibrium.

Thus, we must search for equilibria where the support of player 1's randomized strategy is $\{T, B\}$, and the support of player 2's randomized strategy is $\{L, M, R\}$ or $\{M, R\}$ or $\{L, M\}$ or $\{L, R\}$.

We consider these alternative supports in this order.

Guess support is $\{T, B\}$ for 1 and $\{L, M, R\}$ for 2?

We may denote 1's strategy by $p[T] + (1-p)[B]$ and 2's strategy by $q[L] + (1-q-r)[M] + r[R]$, that is $p = \sigma_1(T)$, $1-p = \sigma_1(B)$, $q = \sigma_2(L)$, $r = \sigma_2(R)$, $1-q-r = \sigma_2(M)$.

Player 1 randomizing over $\{T, B\}$ requires $w_1 = Eu_1(T|\sigma_2) = Eu_1(B|\sigma_2)$,

and so $w_1 = 7q + 2(1-q-r) + 3r = 2q + 7(1-q-r) + 4r$.

Player 2 randomizing over $\{L, M, R\}$ requires $w_2 = Eu_2(L|\sigma_1) = Eu_2(M|\sigma_1) = Eu_2(R|\sigma_1)$,

and so $w_2 = 2p + 7(1-p) = 7p + 2(1-p) = 6p + 5(1-p)$.

We have three equations for three unknowns (p, q, r) , but they have no solution (as the two indifference equations for player 2 imply both $p=1/2$ and $p=3/4$, which is impossible).

Thus there is no equilibrium with this support.

Guess support is $\{T, B\}$ for 1 and $\{M, R\}$ for 2?

We may denote 1's strategy by $p[T] + (1-p)[B]$ and 2's strategy by $(1-r)[M] + r[R]$. ($q=0$)

Player 1 randomizing over $\{T, B\}$ requires $w_1 = Eu_1(T|\sigma_2) = Eu_1(B|\sigma_2)$, so $w_1 = 2(1-r) + 3r = 7(1-r) + 4r$.

Player 2 randomizing over $\{M, R\}$ requires $w_2 = Eu_2(M|\sigma_1) = Eu_2(R|\sigma_1)$, so $w_2 = 7p + 2(1-p) = 6p + 5(1-p)$.

These solution for these two equations in two unknowns is $p=3/4$ and $r=5/4$.

But this solution would yield $\sigma_2(M) = 1-r = -1/4 < 0$, and so there is no equilibrium with this support.

(Notice: if player 2 never chose L then T would be dominated by B for player 1.)

Guess support is $\{T, B\}$ for 1 and $\{L, M\}$ for 2?

We may denote 1's strategy by $p[T] + (1-p)[B]$ and 2's strategy by $q[L] + (1-q)[M]$. ($r=0$)

Player 1 randomizing over $\{T, B\}$ requires $w_1 = Eu_1(T|\sigma_2) = Eu_1(B|\sigma_2)$, so $w_1 = 7q + 2(1-q) = 2q + 7(1-q)$.

Player 2 randomizing over $\{L, M\}$ requires $w_2 = Eu_2(L|\sigma_1) = Eu_2(M|\sigma_1)$, so $w_2 = 2p + 7(1-p) = 7p + 2(1-p)$.

These solution for these two equations in two unknowns is $p=1/2$ and $q=1/2$, with $w_1=4.5 = w_2$.

This solution yields nonnegative probabilities for all actions.

But we also need to check that player 2 would not prefer deviating outside her support to R.

However $Eu_2(R|\sigma_1) = 6p + 5(1-p) = 6 \times 1/2 + 5 \times 1/2 = 5.5 > w_2 = Eu_2(L|\sigma_1) = 2 \times 1/2 + 7 \times 1/2 = 4.5$.

So there is no equilibrium with this support.

Guess support is $\{T, B\}$ for 1 and $\{L, R\}$ for 2?

We may denote 1's strategy by $p[T] + (1-p)[B]$ and 2's strategy by $q[L] + (1-q)[R]$. ($r=1-q$)

Player 1 randomizing over $\{T, B\}$ requires $w_1 = Eu_1(T|\sigma_2) = Eu_1(B|\sigma_2)$, so $w_1 = 7q + 3(1-q) = 2q + 4(1-q)$.

Player 2 randomizing over $\{L, R\}$ requires $w_2 = Eu_2(L|\sigma_1) = Eu_2(R|\sigma_1)$, so $w_2 = 2p + 7(1-p) = 6p + 5(1-p)$.

These solution for these two equations in two unknowns is $p=1/3$ and $q=1/6$.

This solution yields nonnegative probabilities for all actions.

We also need to check that player 2 would not prefer deviating outside her support to M;

$Eu_2(M|\sigma_1) = 7p + 2(1-p) = 7 \times 1/3 + 2 \times 2/3 = 11/3 < w_2 = Eu_2(L|\sigma_1) = 2 \times 1/3 + 7 \times 2/3 = 16/3$.

Thus, we have an equilibrium with this support: $((1/3)[T] + (2/3)[B], (1/6)[L] + (5/6)[R])$.

The expected payoffs in this equilibrium are $w_1 = Eu_1 = 7 \times 1/6 + 3 \times 5/6 = 2 \times 1/6 + 4 \times 5/6 = 11/3 = 3.667$

and $w_2 = Eu_2 = 2 \times 1/3 + 7 \times 2/3 = 6 \times 1/3 + 5 \times 2/3 = 16/3 = 5.333$.

Analysis of the general War of Attrition game.

The game has two given parameters, T and K , where T is the largest number of days that the two players can fight, and K is the value of the prize. (The example in assignment 1 had $T=2$, $K=9$.) There are two players numbered 1 and 2.

Each player i must choose a number a_i in the set $\{0,1,\dots,T\}$.

Here i 's decision a_i represents the number of days that player i is prepared to fight for the prize.

A player wins the prize only if he is prepared to fight strictly longer than the other player.

They will fight for as many days as both are prepared to fight.

Each day of fighting costs each player one dollar, and the prize is worth K dollars.

Assume utility is money, and so the utility payoffs for players 1 and 2 are as follows:

Player 1's payoff is $u_1(a_1,a_2) = K - a_2$ if $a_1 > a_2$, but $u_1(a_1,a_2) = -a_1$ if $a_1 \leq a_2$.

Player 2's payoff is $u_2(a_1,a_2) = K - a_1$ if $a_2 > a_1$, but $u_2(a_1,a_2) = -a_2$ if $a_2 \leq a_1$.

Let us look for a symmetric randomized equilibrium where each player uses a randomized strategy $\sigma = (\sigma(0),\sigma(1),\dots,\sigma(T))$ that assigns positive probability $\sigma(c) > 0$ to every c in $\{0,1,\dots,T\}$.

So player 1 must get the same expected payoff $Eu_1(c,\tilde{a}_2)$ from choosing any pure strategy $a_1=c$ when player 2 uses the randomized strategy σ to randomly determine \tilde{a}_2 .

Notice first that $Eu_1(0,\tilde{a}_2) = 0$.

More generally,

$$Eu_1(c,\tilde{a}_2) = (K-0)\sigma(0) + (K-1)\sigma(1) + \dots + (K-(c-1))\sigma(c-1) - c(\sigma(c) + \dots + \sigma(T)).$$

Compare the results for player 1 of preparing to fight an additional day:

$$Eu_1(c+1,\tilde{a}_2) = (K-0)\sigma(0) + (K-1)\sigma(1) + \dots + (K-c)\sigma(c) - (c+1)(\sigma(c+1) + \dots + \sigma(T)).$$

$$\text{Thus, } Eu_1(c+1,\tilde{a}_2) = Eu_1(c,\tilde{a}_2) + K\sigma(c) - (\sigma(c+1) + \dots + \sigma(T)) = Eu_1(c,\tilde{a}_2) + K\sigma(c) - \sum_{a>c} \sigma(a).$$

That is, being prepared to fight $c+1$ days instead of c could increase 1's payoff by K with probability $\sigma(c)$ but could also decrease 1's payoff by 1 with probability $\sum_{a>c} \sigma(a)$.

It will be helpful to rewrite this equation as by adding and subtracting $\sigma(c)$, to get

$$Eu_1(c+1,\tilde{a}_2) = Eu_1(c,\tilde{a}_2) + (K+1)\sigma(c) - \sum_{a\geq c} \sigma(a).$$

To make player 1 indifferent among all pure strategies, we must have

$$Eu_1(c+1,\tilde{a}_2) = Eu_1(c,\tilde{a}_2) = Eu_1(0,\tilde{a}_2) = 0 \text{ for all } c \text{ in } \{0, 1, 2, \dots, T-1\}.$$

So for all c in $\{0,1,\dots,T-1\}$, we must have

$$0 = Eu_1(c+1,\tilde{a}_2) - Eu_1(c,\tilde{a}_2) = (K+1)\sigma(c) - (\sum_{a\geq c} \sigma(a)).$$

That is, $\sigma(c) = \sum_{a\geq c} \sigma(a)/(K+1)$ for all $c \in \{0,1,\dots,T-1\}$.

But $\sum_{a\geq 0} \sigma(a) = \sigma(0) + \sigma(1) + \dots + \sigma(T) = 1$, and so $\sum_{a\geq c} \sigma(a) = 1 - \sum_{a<c} \sigma(a)$.

Thus, for all c in $\{0,1,\dots,T-1\}$, we have $\sigma(c) = (1 - \sum_{a<c} \sigma(a))/(K+1)$.

This equation can be used to compute $\sigma(0), \sigma(1), \dots, \sigma(T-1), \sigma(T)$:

At $c=0$, we have $\sum_{a<0} \sigma(a) = 0$, and so $\sigma(0) = 1/(K+1)$.

Then $\sigma(1) = (1 - \sigma(0))/(K+1)$, $\sigma(2) = (1 - \sigma(0) - \sigma(1))/(K+1)$, ...

and so on, up to $\sigma(T-1) = (1 - \sigma(0) - \sigma(1) - \dots - \sigma(T-2))/(K+1)$,

Finally the last probability must be $\sigma(T) = 1 - \sigma(0) - \dots - \sigma(T-1)$.

It can be shown that these formulas yield the general solution:

$$\sigma(c) = K^c / (K+1)^{c+1} \text{ for } c = 0,1,\dots,T-1, \text{ and } \sigma(T) = (K / (K+1))^T.$$

As T goes to infinity, this terminal probability $\sigma(T)$ goes to 0.

A finitely repeated game. Consider a game where, in each period, the players play the following game in which each must decide whether to "Fight" or "NotFight," and payoffs are

	f_2	n_2
f_1	-1, -1	9, 0
n_1	0, 9	0, 0

If played only **once**, this game has three equilibria: (f_1, n_2) yielding (9,0), (n_1, f_2) yielding (0, 9), and $(0.9[f_1]+0.1[n_1], 0.9[f_2]+0.1[n_2])$ yielding expected payoffs (0,0).

Suppose this is played **twice**, and period-2 behavior can depend on the outcome in period 1. The overall goal of each player i is to maximize $u_i(1)+\delta u_i(2)$, where $u_i(t)$ is i 's payoff in period t . In a subgame perfect equilibrium of the overall two-period game, the players' anticipated behavior in the final period 2 must look like an equilibrium of the one-period game, given whatever happened in period 1. But the players' understanding of which equilibrium they will play in the second period may depend on the outcome of their play in the first period.

We may say that the state of the players' shared understanding in period 2 will be "state 1" if they expect to play the (f_1, n_2) equilibrium in period 2, "state 2" if they expect to play the (n_1, f_2) equilibrium in period 2, and "state 0" if they expect the randomized equilibrium in period 2.

Consider a subgame-perfect eqm where period-2 depends on period-1 play as follows:

if (f_1, n_2) is played in period 1 then they anticipate state 1 ((f_1, n_2) again) in period 2;

if (n_1, f_2) is played in period 1 then they anticipate state 2 ((n_1, f_2) again) in period 2;

and otherwise they anticipate state 0 (the randomized equilibrium) in period 2.

When the first-period influences second-period behavior in this way, total discounted payoffs for the two players depend on the first-period moves as follows:

	f_2	n_2
f_1	-1+0 δ , -1+0 δ	9+9 δ , 0+0 δ
n_1	0+0 δ , 9+9 δ	0+0 δ , 0+0 δ

(For discounted average value over two periods, we would divide all these payoffs by $1+\delta$.)

So there are three possible equilibria in the first period: (f_1, n_2) yielding total expected payoffs $(9+9\delta, 0)$, (n_1, f_2) yielding total expected payoffs $(0, 9+9\delta)$, and a symmetric randomized equilibrium where each player fights with probability $p = (9+9\delta)/(10+9\delta)$ and each player's expected total payoff is just 0.

But there are also other subgame-perfect equilibria. For example, the anticipated second-period equilibrium might depend on first-period play as follows:

if (f_1, n_2) is played in period 1 then they anticipate state 2 (switch to (n_1, f_2)) in period 2;

if (n_1, f_2) is played in period 1 then they anticipate state 1 (switch to (f_1, n_2)) in period 2;

and otherwise they anticipate state 0 (the randomized equilibrium) in period 2.

Then total discounted payoffs for the two players depend on the first-period moves as follows:

	f_2	n_2
f_1	-1+0 δ , -1+0 δ	9+0 δ , 0+9 δ
n_1	0+9 δ , 9+0 δ	0+0 δ , 0+0 δ

So there are three possible equilibria in period 1: (f_1, n_2) yielding total expected payoffs $(9, 9\delta)$, (n_1, f_2) yielding total expected payoffs $(9\delta, 9)$, and a symmetric randomized equilibrium where each player fights with probability $p = 9/(10 + 9\delta)$ and each player's expected total payoff is $81\delta/(10 + 9\delta)$, which is about 4.26 when δ is close to 1.

Introduction to repeated games Players 1 and 2 will meet on $\tau+1$ days, numbered $0,1,2,\dots,\tau$.

On each day, each player i must choose to be generous (g_i) or selfish (f_i).

On each day k , they get payoffs (u_{1k}, u_{2k}) that depend on their actions (c_{1k}, c_{2k}) as follows:

Player 1: \	Player 2:	g_2	f_2	
g_1		3, 3	0, 5	(Prisoners' dilemma)
f_1		5, 0	2, 2	

except on the last day τ their payoffs will be:

Player 1: \	Player 2:	g_2	f_2	
g_1		5, 5	0, 4	(Trust game)
f_1		4, 0	2, 2	

On each day, each player knows what both players did on all previous days.

Each player wants to maximize the expected discounted sum of his payoffs $V_i = u_{i0} + \delta u_{i1} + \delta^2 u_{i2} + \dots + \delta^\tau u_{i\tau}$ for some given discount factor δ between 0 and 1.

If the Prisoners' Dilemma were played once, (f_1, f_2) would be the only equilibrium, yielding the Pareto-dominated payoffs (2,2). And for this multi-period game, both players doing f_i always is one equilibrium. But in multi-period games, opportunities to respond later can enlarge the set of equilibria.

Consider the strategy for each player i to **choose g_i until f_1 or f_2 is chosen, but thereafter choose f_i .**

We can show that, if $\delta \geq 2/3$, it is a subgame-perfect equilibrium for both players to choose this strategy.

Consider first the case of $\tau=1$, where the prisoners' dilemma is played once, followed by one play of the trust game at the end. Under the strategies described here, on the last day,

they will play the good (g_1, g_2) equilibrium of the "trust game" if both were previously generous, but they will play the bad (f_1, f_2) equilibrium if either player was previously selfish.

So the overall payoffs will depend on their first-day choices as follows:

Player 1: \	Player 2:	g_2	f_2
g_1		$3 + \delta 5, 3 + \delta 5$	$0 + \delta 2, 5 + \delta 2$
f_1		$5 + \delta 2, 0 + \delta 2$	$2 + \delta 2, 2 + \delta 2$

Then (g_1, g_2) is an equilibrium at the first day if $3 + 5\delta \geq 5 + 2\delta$, that is, if $\delta \geq 2/3$.

A similar calculation can be made for any number $\tau \geq 1$ of repetitions of the prisoners' dilemma.

Let $G(\tau)$ be the discounted sum of payoffs from (g_1, g_2) -always, and let $F(\tau)$ be the discounted sum of payoffs from (f_1, f_2) always, in τ repetitions of the prisoners' dilemma followed by one trust game.

So $G(0) = 5$ and $F(0) = 2$ and, for any $\tau \geq 1$, $G(\tau) = 3 + \delta G(\tau-1)$ and $F(\tau) = 2 + \delta F(\tau-1)$.

Fact: $w + w\delta + w\delta^2 + \dots + w\delta^{s-1} = w(1 - \delta^s)/(1 - \delta)$. So $G(\tau) = 3(1 - \delta^\tau)/(1 - \delta) + 5\delta^\tau$, $F(\tau) = 2(1 - \delta^{\tau+1})/(1 - \delta)$.

Lemma: If $1 > \delta \geq 2/3$ then $G(\tau) - F(\tau) \geq 3$ for all τ . (Proof by induction: $G(0) - F(0) = 5 - 2 = 3$, and then for any $\tau \geq 1$ we get inductively $G(\tau) - F(\tau) = 3 - 2 + \delta(G(\tau-1) - F(\tau-1)) \geq 1 + (2/3)(3) = 3$.)

Now assuming that the strategies described above will be played after the first stage, the players' overall payoffs will depend on their first-day choices as follows:

Player 1: \	Player 2:	g_2	f_2
g_1		$3 + \delta G(\tau-1), 3 + \delta G(\tau-1)$	$0 + \delta F(\tau-1), 5 + \delta F(\tau-1)$
f_1		$5 + \delta F(\tau-1), 0 + \delta F(\tau-1)$	$2 + \delta F(\tau-1), 2 + \delta F(\tau-1)$

With $1 > \delta \geq 2/3$, for any τ , it is an equilibrium for both to start doing g_i , as these strategies specify, because $3 + \delta G(\tau-1) \geq 5 + \delta F(\tau-1)$. (*Proof:* $3 + \delta G(\tau-1) - (5 + \delta F(\tau-1)) = -2 + \delta(G(\tau-1) - F(\tau-1)) \geq -2 + (2/3)(3) = 0$.)

As $\tau \rightarrow \infty$, $G(\tau) \rightarrow 3/(1 - \delta)$, $F(\tau) \rightarrow 2/(1 - \delta)$, and overall payoffs here depend on first-day actions as follows:

Player 1: \	Player 2:	g_2	f_2
g_1		$3 + \delta 3/(1 - \delta), 3 + \delta 3/(1 - \delta)$	$0 + \delta 2/(1 - \delta), 5 + \delta 2/(1 - \delta)$
f_1		$5 + \delta 2/(1 - \delta), 0 + \delta 2/(1 - \delta)$	$2 + \delta 2/(1 - \delta), 2 + \delta 2/(1 - \delta)$

The equilibrium condition $3 + \delta 3/(1 - \delta) \geq 5 + \delta 2/(1 - \delta)$ is satisfied when $1 > \delta \geq 2/3$.

Repeated games Infinitely repeated games are useful models of long-term relationships.

The game will be played at an infinite sequence of time periods numbered 1,2,3,...

Suppose that the set of players is $\{1,2\}$. In each period k , each player i must choose an action a_{ik} in some set A_i . In period k , each player i 's payoff u_{ik} will depend on both players' actions according to some utility function $U_i: A_1 \times A_2 \rightarrow \mathbb{R}$; that is $u_{ik} = U_i(a_{1k}, a_{2k})$.

We assume here that the actions at each period are publicly observable, and so each player's action in each period may depend on the history of actions by both players at all past periods.

Given any discount factor δ such that $0 \leq \delta < 1$, the δ -discounted average value of player i 's payoffs is $DAV(u_{i1}, u_{i2}, u_{i3}, \dots) = (1-\delta)(u_{i1} + \delta u_{i2} + \delta^2 u_{i3} + \dots + \delta^{k-1} u_{ik} + \dots)$.

(For any x , $DAV(x, x, x, \dots) = x$. If δ is slightly less than 1 then the players are very patient.)

The objective of each player i in the repeated game is to maximize the expected discounted average value of his payoffs, with respect to some discount factor δ , where $0 < \delta < 1$.

Fact. (Recursion formula) $DAV(u_{i1}, u_{i2}, u_{i3}, \dots) = (1-\delta)u_{i1} + \delta DAV(u_{i2}, u_{i3}, u_{i4}, \dots)$.

We may describe equilibria of repeated games in terms of a various social states.

At each period of the game, the players will understand that their current relationship is described by one of these social states, and their expectations about each others' behavior will be determined by this state. This state may be called the state of play in the game at this period.

(These social states are an attribute of the equilibrium, not of the game, as they describe the different kinds of expectations that the players may have about each others' future behavior.)

To describe an equilibrium or scenario in terms of social states, we must specify the following:

(1) Social states We must list the set of social states in this equilibrium. (States may denoted by numbers or may be named for the kinds of interpersonal relationships that they represent.)

(2) State-dependent strategies. For each state θ , we must specify a profile of (possibly randomized) actions $(\tilde{s}_1(\theta), \tilde{s}_2(\theta))$ describing the predicted behavior of the players in any period when this θ is the state of play.

(3) Transitions. For each social state θ , we must specify the profiles of players' actions that would cause the state of play in the next period to change from this state to another state. We may let $\Theta(a_1, a_2; \theta)$ denote the state of play in the next period after a period when the state of play was θ and the players chose actions (a_1, a_2) (possibly deviating from the prediction $(\tilde{s}_1(\theta), \tilde{s}_2(\theta))$).

(4) Initial state. We must specify which social state is initial state of play in the first period of the game. Here we will generally let state "0" denote this initial state.

Given any scenario as in (1)-(3) above, and given any discount factor δ , let $V_i(\theta)$ denote the expected δ -discounted average value of player i 's payoffs in this scenario when (ignoring (4)) the state of play begins in state θ . Given $\delta < 1$, these numbers $V_i(\theta)$ can be computed (with algebra) from the equations: $V_i(\theta) = E[(1-\delta)U_i(\tilde{s}_1(\theta), \tilde{s}_2(\theta))] + \delta V_i(\Theta(\tilde{s}_1(\theta), \tilde{s}_2(\theta); \theta))$.

Fact. A scenario as in (1)-(3) above is a subgame-perfect equilibrium if, for every player i and every state θ , player i could not expect to gain by unilaterally deviating from the prediction $\tilde{s}_i(\theta)$ in a period when the state of play is θ . That is, we have an equilibrium if, for every state θ ,

$V_1(\theta) \geq E[(1-\delta)U_1(a_1, \tilde{s}_2(\theta))] + \delta V_1(\Theta(a_1, \tilde{s}_2(\theta); \theta))$, for all a_1 in A_1 ,

$V_2(\theta) \geq E[(1-\delta)U_2(\tilde{s}_1(\theta), a_2)] + \delta V_2(\Theta(\tilde{s}_1(\theta), a_2; \theta))$, for all a_2 in A_2 .

(This is the "one-deviation property" of Osborne p 438.)

Example 1. Consider a repeated game where, in each period, the players play the following "Prisoners' dilemma" game in which each must decide whether to "cooperate" or "defect".

	c_2	d_2
c_1	3, 3	0, 5
d_1	5, 0	2, 2

We first consider a version of the "grim trigger" equilibrium:

The states are $\{0, 1\}$. (State 0 represents "trust" or "friendship"; state 1 represents "distrust".)

The predicted behavior in state 0 is (c_1, c_2) . The predicted behavior in state 1 is (d_1, d_2) .

In any period when the current state of play is 0, if the players' action profile is (c_1, d_2) or (d_1, c_2) then the state of play next period will switch to state 1, otherwise it will remain state 0.

When the state of play is 1, the future state of play always remains state 1.

The expected discounted average values for the players in the states satisfy the equations:

$$V_1(0) = (1-\delta)U_1(c_1, c_2) + \delta V_1(0), \quad V_1(1) = (1-\delta)U_1(d_1, d_2) + \delta V_1(1),$$

$$V_2(0) = (1-\delta)U_2(c_1, c_2) + \delta V_2(0), \quad V_2(1) = (1-\delta)U_2(d_1, d_2) + \delta V_2(1).$$

$$\text{So } V_1(0) = (1-\delta)3 + \delta V_1(0), \quad V_1(1) = (1-\delta)2 + \delta V_1(1), \quad \text{and so } V_1(0) = 3, \quad V_1(1) = 2.$$

$$\text{Similarly, } V_2(0) = 3, \quad V_2(1) = 2.$$

For this scenario to be an equilibrium, we need:

$$V_1(0) \geq (1-\delta)U_1(d_1, c_2) + \delta V_1(1), \quad V_1(1) \geq (1-\delta)U_1(c_1, d_2) + \delta V_1(1),$$

$$V_2(0) \geq (1-\delta)U_2(c_1, d_2) + \delta V_2(1), \quad V_2(1) \geq (1-\delta)U_2(d_1, c_2) + \delta V_2(1).$$

That is, we need: $3 \geq (1-\delta)5 + \delta 2$ and $2 \geq (1-\delta)0 + \delta 2$, which are satisfied when $1 \geq \delta \geq 2/3$.

Now let's consider another (more forgiving) equilibrium:

The states are $\{0, 1, 2\}$. (state 0 is "friendship"; state 1 is "punishing 1"; state 2 is "punishing 2".)

The predicted behavior in state 0 is (c_1, c_2) . The predicted behavior in state 1 is (c_1, d_2) .

The predicted behavior in state 2 is (d_1, c_2) .

When the state of play is 0, if the players choose (d_1, c_2) then the next state of play will be 1, if the players choose (c_1, d_2) then the state of play next period will be 2, otherwise it will remain 0.

When the state of play is 1, if the players choose (c_1, d_2) then the next state of play will be 0, otherwise it will remain 1. When the state of play is 2, if the players choose (d_1, c_2) then the next state of play will be 0, otherwise it will remain 2.

The expected discounted average values $V_i(\theta)$ for player 1 in each state θ satisfy the equations:

$$V_1(0) = (1-\delta)U_1(c_1, c_2) + \delta V_1(0), \quad V_1(1) = (1-\delta)U_1(c_1, d_2) + \delta V_1(0),$$

$$V_1(2) = (1-\delta)U_1(d_1, c_2) + \delta V_1(0).$$

$$\text{So } V_1(0) = (1-\delta)3 + \delta V_1(0), \quad \text{and } V_1(0) = 3. \quad \text{So } V_1(1) = (1-\delta)0 + \delta 3 = 3\delta,$$

$$\text{and } V_1(2) = (1-\delta)5 + \delta 3 = 5 - 2\delta.$$

For a subgame-perfect equilibrium, the values for player 1 must satisfy:

$$V_1(0) \geq (1-\delta)U_1(d_1, c_2) + \delta V_1(1), \quad V_1(1) \geq (1-\delta)U_1(d_1, d_2) + \delta V_1(1),$$

$$V_1(2) \geq (1-\delta)U_1(c_1, c_2) + \delta V_1(2).$$

$$\text{These inequalities become respectively: } 3 \geq (1-\delta)5 + \delta 3\delta, \quad 3\delta \geq (1-\delta)2 + \delta 3\delta,$$

$$5 - 2\delta \geq (1-\delta)3 + \delta(5 - 2\delta); \quad \text{and these are all satisfied when } 1 \geq \delta \geq 2/3.$$

(Algebraic fact used here: $(1-\delta^2) = (1-\delta)(1+\delta)$.)

The values equations for player 2 are similarly: $V_2(0) = (1-\delta)U_2(c_1, c_2) + \delta V_2(0),$

$$V_2(1) = (1-\delta)U_2(c_1, d_2) + \delta V_2(0), \quad V_2(2) = (1-\delta)U_2(d_1, c_2) + \delta V_2(0);$$

$$\text{and these equations imply } V_2(0) = 3, \quad V_2(1) = 5 - 2\delta, \quad V_2(2) = 3\delta.$$

Then the equilibrium inequalities for player 2 are: $V_2(0) \geq (1-\delta)U_2(c_1, d_2) + \delta V_2(2),$

$$V_2(1) \geq (1-\delta)U_2(c_1, c_2) + \delta V_2(1), \quad V_2(2) \geq (1-\delta)U_2(d_1, d_2) + \delta V_2(2).$$

$$\text{These inequalities respectively become } 3 \geq (1-\delta)5 + \delta 3\delta, \quad 5 - 2\delta \geq (1-\delta)3 + \delta(5 - 2\delta), \quad 3\delta \geq (1-\delta)2 + \delta 3\delta;$$

and these are also satisfied when $1 \geq \delta \geq 2/3$.

Example 2 Consider a repeated game where players 1 and 2 repeated play the game below infinitely often. In each round, each player i must decide whether to fight (f_i) or not (n_i). The players want to maximize their δ -discount average value of payoffs, for some $0 < \delta < 1$.

	f_2	n_2
f_1	-1, -1	9, 0
n_1	0, 9	0, 0

A subgame-perfect equilibrium:

States: there are three states, numbered 0,1,2. The initial state in period 1 is state 0.

(State 1 may be interpreted as "1 has ownership", state 2 may be interpreted as "2 has ownership" and state 0 may be interpreted as "fighting for ownership" or *war of attrition*.)

Strategies: Let $s_i(\theta)$ denote the move that player i would choose in state θ .

Player 1's strategy is $s_1(1) = f_1$, $s_1(2) = n_1$, $s_1(0) = q[f_1] + (1-q)[n_1]$ for some q between 0 and 1.

Player 2's strategy is $s_2(1) = n_1$, $s_2(2) = f_1$, $s_2(0) = q[f_2] + (1-q)[n_2]$ for the same q .

We will need to find what q makes this an equilibrium.

Transitions: When the current state is state 0, the state next period would be:

state 1 if (f_1, n_2) is played now, state 2 if (n_1, f_2) is played now, and state 0 if (f_1, f_2) or (n_1, n_2) is played now. Once the game is in state 1 or 2, it stays in the same state forever.

Values: Let $V_i(\theta)$ denote the expected discounted average value of payoffs for player i in state θ .

The recursion equations for states 1 and 2 are

$$V_i(1) = (1-\delta)U_i(f_1, n_2) + \delta V_i(1), \text{ for } i=1,2, \text{ and so } V_1(1) = 9 \text{ and } V_2(1) = 0;$$

$$V_i(2) = (1-\delta)U_i(n_1, f_2) + \delta V_i(2), \text{ for } i=1,2, \text{ and so } V_1(2) = 0 \text{ and } V_2(2) = 9.$$

To check the equilibrium condition in state 1, notice that

$$9 = V_1(1) \geq (1-\delta)U_1(n_1, n_2) + \delta V_1(1) = (1-\delta)(0) + \delta(9) = \delta 9,$$

$$0 = V_2(1) \geq (1-\delta)U_2(f_1, f_2) + \delta V_2(1) = (1-\delta)(-1) + \delta(0) = -(1-\delta).$$

The equilibrium conditions in state 2 are similar.

In state 0, for player 1 to be willing to randomize between f_1 and n_1 , he must expect the same discounted average value $V_1(0)$ from choosing f_1 or n_1 this period, and so we must have

$$V_1(0) = q((1-\delta)U_1(f_1, f_2) + \delta V_1(0)) + (1-q)((1-\delta)U_1(f_1, n_2) + \delta V_1(1)), \text{ and}$$

$$V_1(0) = q((1-\delta)U_1(n_1, f_2) + \delta V_1(2)) + (1-q)((1-\delta)U_1(n_1, n_2) + \delta V_1(0)).$$

The latter is $V_1(0) = q(1-\delta)0 + q\delta 0 + (1-q)(1-\delta)0 + (1-q)\delta V_1(0)$, implying $V_1(0) = 0$.

Then $V_1(0) = q(1-\delta)(-1) + q\delta V_1(0) + (1-q)(1-\delta)9 + (1-q)\delta 9$, implies $q = 9/(10-\delta)$.

Similarly, $V_2(0) = 0$.

Another subgame-perfect equilibrium for this repeated game, with δ -discounting.

	f_2	n_2
f_1	-1, -1	9, 0
n_1	0, 9	0, 0

States: there are three states, numbered 0,1,2. The initial state in period 1 is state 0.

(In this equilibrium, state 1 may be interpreted as "1's turn", state 2 may be interpreted as "2's turn" and state 0 may be interpreted as "confused about whose turn it is".)

Strategies: Let $s_i(\theta)$ denote the move that player i would choose in state θ .

Player 1's strategy is $s_1(1) = f_1$, $s_1(2) = n_1$, $s_1(0) = q[f_1] + (1-q)[n_1]$ for some q between 0 and 1.

Player 2's strategy is $s_2(1) = n_2$, $s_2(2) = f_2$, $s_2(0) = q[f_2] + (1-q)[n_2]$ for the same q as player 1.

We will need to find what q makes this an equilibrium.

Transitions: When the current state is state 0, the state next period would be: state 2 if (f_1, n_2) is played now, state 1 if (n_1, f_2) is played now, and state 0 if (f_1, f_2) or (n_1, n_2) is played now.

When the current state is 1, the next state is always 2. When the current state is 2, the next state is always 1. (So from state 1 or 2, the state of play alternates between states 1 and 2 forever.)

Values Let $V_i(\theta)$ denote the expected discounted average value of payoffs for player i in state θ .

The recursion equations for states 1 and 2 are

$$V_i(1) = (1-\delta)U_i(f_1, n_2) + \delta V_i(2) \text{ and } V_i(2) = (1-\delta)U_i(n_1, f_2) + \delta V_i(1) \text{ for } i=1,2.$$

$$\text{So } V_1(1) = (1-\delta)9 + \delta[(1-\delta)0 + \delta V_1(1)], \text{ and so } V_1(1) = 9(1-\delta)/(1-\delta^2) = 9/(1+\delta)$$

$$\text{and } V_1(2) = (1-\delta)0 + \delta V_1(1) = \delta 9/(1+\delta).$$

$$\text{Similarly, } V_2(2) = 9/(1+\delta) \text{ and } V_2(1) = \delta 9/(1+\delta).$$

To check the equilibrium condition in state 1, notice that

$$9/(1+\delta) = V_1(1) \geq (1-\delta)U_1(n_1, n_2) + \delta V_1(2) = \delta^2 9/(1+\delta),$$

$$\delta 9/(1+\delta) = V_2(1) \geq (1-\delta)U_2(f_1, f_2) + \delta V_2(2) = (1-\delta)(-1) + \delta 9/(1+\delta).$$

The equilibrium conditions in state 2 are similar.

In state 0, for player 1 to be willing to randomize between f_1 and n_1 , he must expect the same discounted average value $V_1(0)$ from choosing f_1 or n_1 this period, and so we must have

$$V_1(0) = q((1-\delta)U_1(f_1, f_2) + \delta V_1(0)) + (1-q)((1-\delta)U_1(f_1, n_2) + \delta V_1(2)), \text{ and}$$

$$V_1(0) = q((1-\delta)U_1(n_1, f_2) + \delta V_1(1)) + (1-q)((1-\delta)U_1(n_1, n_2) + \delta V_1(0)).$$

$$\text{So } V_1(0) = q(1-\delta)(-1) + q\delta V_1(0) + (1-q)(1-\delta)9 + (1-q)\delta^2 9/(1+\delta),$$

$$\text{and } V_1(0) = q(1-\delta)0 + q\delta 9/(1+\delta) + (1-q)(1-\delta)0 + (1-q)\delta V_1(0).$$

These two equations can be solved for the two unknowns $V_1(0)$ and q .

The explicit formula is hard to derive, but the equations can be solved numerically on a computer. The second equation implies $V_1(0) = [q\delta 9/(1+\delta)]/[1-\delta(1-q)]$, and substituting this into the first equation yields a nonlinear equation in one unknown q , which can be solved with Excel's Goal-Seek tool. The results with $\delta = 0.99$ are $V_1(0) = 4.446$ and $q = 0.585$

The value for player 2 in state 0 is of course the same, $V_2(0) = V_1(0)$, because everything is symmetric as long as they are in state 0.

A First Bayesian Game Bayesian games are models of one-stage games where players choose actions simultaneously, but where each player may have private information, called his type. Let us consider an example where player 2 is uncertain about one of player 1's payoffs. Each player must independently decide whether to act with friendship (f_i) or aggression (g_i). Player 1 might be the kind of person who would be contented (type 1c) or envious (type 1e) if they chose to be friendly. Player 2 thinks that each of 1's possible types has probability 0.5. The players' payoffs (u_1, u_2) depend on their actions and 1's type as follows:

If 1's type is 1c:	f_2	g_2	$p(1c) = 0.5$
f_1	8,8	0,6	
g_1	6,0	3,3	
If 1's type is 1e:	f_2	g_2	$p(1e) = 0.5$
f_1	5,8	0,6	
g_1	6,0	3,3	

How shall we analyze about this game? Let me first sketch a common mistake. To deal with the uncertainty about 1's payoff from (f_1, f_2), some students try to analyze the game where player 1's payoff from (f_1, f_2) is the expected utility $0.5(8)+0.5(5)=6.5$. So these students consider a 2×2 payoff matrix that differs from the second (1e) case only in that the payoff 5 would be replaced by 6.5, and then they find an "equilibrium" at (f_1, f_2) (as $6.5 > 6$ for player 1 and $8 > 6$ for player 2). Such analysis would be nonsense, however. This "equilibrium" would correspond to a theory that each player is sure to choose friendship. But player 2 knows that if player 1 is type 1e then he will not choose f_1 , because f_1 would be dominated by g_1 for player 1 when his type is 1e. Thus, player 2 must believe that there is at least a probability 0.5 of player 1 being the envious type 1e and thus choosing aggression g_1 . A correct analysis must recognize this fact.

To find a correct approach, we may consider the situation before the players learns any private information, but when they know that each will learn his private type information before he acts in the game. A strategy for a player is a complete plan that specifies a feasible action for the player in every possible contingency that the player could find. Before player 1 learned his type, he would have 4 strategies $\{f_c f_e, f_c g_e, g_c f_e, g_c g_e\}$ because he will learn his type before acting. (For example, $f_c g_e$ denotes the strategy "be friendly if type 1c, be aggressive if type 1e.") Player 2 would have only two strategies $\{f_2, g_2\}$, because she must act without learning 1's type. For each pair of strategies, we can compute the expected payoffs to each player, given that each of 1's types has probability 1/2. So the normal representation in strategic form of this Bayesian game is

	f_2	g_2
$f_c f_e$	6.5, 8	0, 6
$f_c g_e$	7, 4	1.5, 4.5
$g_c f_e$	5.5, 4	1.5, 4.5
$g_c g_e$	6, 0	3, 3

This strategic game has one equilibrium: ($g_c g_e, g_2$), where both are aggressive and get payoffs (3,3). In this strategic game, $f_c f_e$ and $g_c f_e$ are strictly dominated for 1 (by $f_c g_e$ and $g_c g_e$ respectively). When we eliminate these dominated strategies, then f_2 becomes dominated for 2, and $g_c g_e$ is the unique best response for 1 against 2's remaining strategy g_2 . (The students' mistake above was to consider only the strategies $f_c f_e$ and $g_c g_e$ here.)

A Bayesian game is defined by a set of players N ; a set of actions A_i , a set of types T_i , and a utility function $u_i: (\times_{j \in N} A_j) \times (\times_{j \in T} T_j) \rightarrow \mathbb{R}$, for each i in N ; and a probability distribution $p \in \Delta(\times_{j \in N} T_j)$.

Increasing differences and increasing strategies in Bayesian games

We may consider Bayesian games where each player i first learns his type \tilde{t}_i , and then each player i chooses his action a_i . We assume here that each player i 's type is drawn from some probability distribution p_i , independently of all other players' types, and so the joint probability distribution of the players' types can be written $p((t_i)_{i \in N}) = \prod_{i \in N} p_i(t_i)$, where $p_i(t_i) = \text{Prob}(\tilde{t}_i = t_i)$.

The payoffs of each player i in $N = \{1, 2, \dots, n\}$ may depend on all players' types and actions according to some utility payoff function $u_i(c_1, \dots, c_n, \tilde{t}_1, \dots, \tilde{t}_n)$.

A function $f: \mathbb{R} \rightarrow \mathbb{R}$ is increasing (in the weak sense) iff, for all x and \hat{x} , $\hat{x} \geq x$ implies $f(\hat{x}) \geq f(x)$.

A function $f: \mathbb{R} \rightarrow \mathbb{R}$ is strictly increasing iff, for all x and \hat{x} , $\hat{x} > x$ implies $f(\hat{x}) > f(x)$.

Consider a two-player Bayesian game where player 1 has two possible actions, \hat{T} and B .

Player 1 has several possible types, and each possible type is represented by a **number** t_1 .

Player 2 may have many possible actions c_2 and many possible types t_2 .

Suppose that player 2's type t_2 is independent of player 1's type t_1 .

The difference in player 1's payoff in switching from B to T is $u_1(\hat{T}, c_2, t_1, t_2) - u_1(B, c_2, t_1, t_2)$.

This difference depends on player 1's type t_1 , player 2's action c_2 , and player 2's type t_2 .

We say that player 1's payoffs satisfy (weakly or strictly) increasing differences if this difference $u_1(\hat{T}, c_2, t_1, t_2) - u_1(B, c_2, t_1, t_2)$ is a (weakly or strictly) increasing function of t_1 , no matter what player 2's action c_2 and type t_2 may be.

That is, increasing differences (in the weak sense) means that, for every r_1 , t_1 , c_2 , and t_2 :

if $r_1 \geq t_1$ then $u_1(\hat{T}, c_2, r_1, t_2) - u_1(B, c_2, r_1, t_2) \geq u_1(\hat{T}, c_2, t_1, t_2) - u_1(B, c_2, t_1, t_2)$.

Strictly increasing differences means that, for every r_1 , t_1 , c_2 , and t_2 :

if $r_1 > t_1$ then $u_1(\hat{T}, c_2, r_1, t_2) - u_1(B, c_2, r_1, t_2) > u_1(\hat{T}, c_2, t_1, t_2) - u_1(B, c_2, t_1, t_2)$.

With increasing differences, 1's higher types find T relatively more attractive than lower types do.

Player 1 is using a cutoff strategy if there is some number θ (the cutoff) such that, for each possible type t_1 of player 1: if $t_1 > \theta$ then type t_1 would choose $[T]$ for sure in this strategy,

if $t_1 < \theta$ then type t_1 would choose $[B]$ for sure in this strategy,

if $t_1 = \theta$ then type t_1 may choose T or B or may randomize in this strategy.

Comparing cutoff strategies, the probability of 1 choosing T decreases as the cutoff θ increases.

Fact. If player 1's payoffs satisfy increasing differences then, no matter what strategy player 2 may use, player 1 will always want to use a cutoff strategy. Thus, when we are looking for equilibria, the increasing-differences property assures us that player 1 must be using a cutoff strategy.

More generally, in games where player 1's action can be any **number** in some range, we say that player 1's payoffs satisfy (weakly or strictly) increasing differences if, for every pair of possible actions c_1 and d_1 such that $c_1 > d_1$, the difference $u_1(c_1, c_2, t_1, t_2) - u_1(d_1, c_2, t_1, t_2)$ is a (weakly or strictly) increasing function of player 1's type t_1 , no matter what player 2's action c_2 and type t_2 may be. (If u_1 is differentiable then $\partial^2 u_1 / \partial c_1 \partial t_1 \geq 0$.)

Fact. If 1's payoffs satisfy increasing differences, then, against any strategy σ_2 of player 2, player 1 will have some best-response strategy $s_1: T_1 \rightarrow A_1$ that is weakly increasing ($r_1 \geq t_1 \Rightarrow s_1(r_1) \geq s_1(t_1)$).

When 1's payoffs have strictly increasing differences then all player 1's best-response strategies must be weakly increasing: if $r_1 > t_1$ and, against some strategy σ_2 for player 2, action c_1 is optimal for type t_1 and action d_1 is optimal for type r_1 , then $d_1 \geq c_1$.

(By optimality, $Eu_1(c_1, \sigma_2, t_1, \tilde{t}_2) - Eu_1(d_1, \sigma_2, t_1, \tilde{t}_2) \geq 0$ but $0 \geq Eu_1(c_1, \sigma_2, r_1, \tilde{t}_2) - Eu_1(d_1, \sigma_2, r_1, \tilde{t}_2)$, which would contradict strictly increasing differences if we had $c_1 > d_1$.)

Example: Player 1's types possible are $\{0, .1, .2, .3\}$, each with probability $1/4$.
 Player 2 has no private information. 1's actions are $\{\hat{T}, B\}$, 2's actions are $\{L, R\}$.
 Given 1's type t_1 , the payoffs (u_1, u_2) are:

	L	R	
\hat{T}	$t_1, 0$	$t_1, -1$	<i>(we use "\hat{T}" for top here because "T" denotes types)</i>
B	$1, 0$	$-1, 3$	

So 1's utility difference in switching from B to T depends on 2's action and 1's type as follows:
 $u_1(\hat{T}, L, t_1) - u_1(B, L, t_1) = t_1 - 1$, $u_1(\hat{T}, R, t_1) - u_1(B, R, t_1) = t_1 + 1$. Notice that these differences increase in t_1 .
 So higher types t_1 always find \hat{T} relatively more attractive than lower types, and player 1 will use a cutoff strategy of the form: "do \hat{T} if $t_1 > \theta$, do B if $t_1 < \theta$, may randomize if $t_1 = \theta$," for some given cutoff value θ .
 Thus, although player 1 has $2^4 = 16$ pure strategies in this Bayesian game (and $2^{16} - 1 = 65535$), we only need to consider 1's cutoff strategies with the following 9 possible supports:

- $(\theta > .3)$: every type would choose [B], so 2 thinks the probability of \hat{T} is $P(\hat{T}) = 0$;
- $(\theta = .3)$: $\{0, .1, .2\}$ would choose [B], but .3 would randomize in some way, so 2 thinks $0 \leq P(\hat{T}) \leq 1/4$;
- $(.2 < \theta < .3)$: $\{0, .1, .2\}$ would choose [B], but .3 would choose $[\hat{T}]$, so 2 thinks $P(\hat{T}) = 1/4$;
- $(\theta = .2)$: $\{0, .1\}$ would choose [B], .2 could randomize, .3 would choose $[\hat{T}]$, so 2 thinks $1/4 \leq P(\hat{T}) \leq 1/2$;
- $(.1 < \theta < .2)$: $\{0, .1\}$ would choose [B], $\{.2, .3\}$ would choose $[\hat{T}]$, so 2 thinks $P(\hat{T}) = 1/2$;
- $(\theta = .1)$: 0 would choose [B], .1 could randomize, $\{.2, .3\}$ would choose $[\hat{T}]$, so 2 thinks $1/2 \leq P(\hat{T}) \leq 3/4$;
- $(0 < \theta < .1)$: 0 would choose [B], $\{.1, .2, .3\}$ would choose $[\hat{T}]$, so 2 thinks $P(\hat{T}) = 3/4$;
- $(\theta = 0)$: 0 could randomize, $\{.1, .2, .3\}$ would choose $[\hat{T}]$, so 2 thinks $3/4 \leq P(\hat{T}) \leq 1$;
- $(\theta < 0)$: every type would choose $[\hat{T}]$, so 2 thinks $P(\hat{T}) = 1$.

If player 2 uses $\sigma_2 = q[L] + (1-q)[R]$, then player 1's optimal cutoff θ would have the property:
 $t_1 \geq \theta \iff qt_1 + (1-q)t_1 = U_1(\hat{T}, \sigma_2, t_1) \geq U_1(B, \sigma_2, t_1) = q(1) + (1-q)(-1)$.
 This implies $q\theta + (1-q)\theta = q(1) + (1-q)(-1)$. So the cutoff θ is optimal for 1 when $q = (\theta + 1)/2$.

There is obviously no equilibrium in which player 2 chooses L for sure or R for sure. (check!)

To make player 2 willing to randomize, we must have $EU_2(L) = EU_2(R)$, that is,

$$P(\hat{T})(0) + (1 - P(\hat{T}))(0) = P(\hat{T})(-1) + (1 - P(\hat{T}))(3), \text{ and so } P(\hat{T}) = 3/4.$$

Here $P(\hat{T})$ denotes the (unconditional) probability of player 1 choosing \hat{T} as assessed by player 2, who does not know 1's type t_1 . But 1's equilibrium strategy σ_1 must specify, for each possible type t_1 in $\{0, .1, .2, .3\}$, the conditional probability $\sigma_1(\hat{T}|t_1)$ of player 1 doing \hat{T} when his type is t_1 .

These unconditional and conditional probabilities of \hat{T} must satisfy the equation: $P(\hat{T}) = \sum_{t_1} p_1(t_1)\sigma_1(\hat{T}|t_1)$.

For a cutoff strategy with $\sigma_1(\hat{T}|t_1) = 1$ for $t_1 > \theta$ and $\sigma_1(\hat{T}|t_1) = 0$ for $t_1 < \theta$, this is $P(\hat{T}) = p_1(\theta)\sigma_1(\hat{T}|\theta) + \sum_{t_1 > \theta} p_1(t_1)$.
 So to get $P(\hat{T}) = 3/4$, the cutoff θ must be between 0 and .1 (0 would choose [B], $\{.1, .2, .3\}$ would choose $[\hat{T}]$).

Now let q denote the probability of 2 choosing L. To make 1's cutoff strategy optimal for him, 2's randomized strategy $q[L] + (1-q)[R]$ must make player 1 prefer B when $t_1 = 0$, but must make player 1 prefer \hat{T} when $t_1 = .1$.

$EU_1(\hat{T}|t_1=0) \leq EU_1(B|t_1=0)$ implies $(q)(0) + (1-q)(0) \leq (q)(1) + (1-q)(-1)$, and so $1/2 \leq q$.

$EU_1(\hat{T}|t_1=.1) \geq EU_1(B|t_1=.1)$ implies $(q)(.1) + (1-q)(.1) \geq (q)(1) + (1-q)(-1)$, and so $q \leq 11/20$.

That is, to get a cutoff θ such that $0 \leq \theta \leq .1$, we must have $1/2 \leq q = (\theta + 1)/2 \leq 11/20$.

So in equilibrium, 1 chooses B if $t_1 = 0$, 1 chooses \hat{T} if $t_1 \geq .1$, and 2 randomizes, choosing L with some probability q that is between $1/2$ and $11/20$.

Now suppose instead player 1 has five possible types $\{0, .1, .2, .3, .4\}$, each with probability $p_1(t_1) = 1/5$.

To make player 2 willing to randomize, player 1 must use a strategy such that $P(\hat{T}) = 3/4$.

For that to occur in an increasing cutoff strategy, the cutoff must be at $\theta = .1$.

So $t_1 = 0$ chooses B; and when $t_1 > .1$ (which has probability $3/5$) player 1 chooses \hat{T} .

The remaining $3/4 - 3/5 = 0.15$ probability of T must come from player 1 choosing \hat{T} with probability $\sigma_1(\hat{T}|.1) = 0.15/p_1(.1) = 0.15/0.2 = 0.75$ when $t_1 = .1$.

To make type $t_1 = \theta = .1$ willing to randomize, 2's probability of choosing L must be $q = (.1 + 1)/2 = 11/20$.

Example. Player 1's type t_1 is drawn from a Uniform distribution on the interval from 0 to 1, and payoffs (u_1, u_2) depend on 1's type as follows, where ε is a given number between 0 and 1 (say $\varepsilon=0.1$):

	L	R
\hat{T}	$\varepsilon t_1, 0$	$\varepsilon t_1, -1$
B	$1, 0$	$-1, 3$

Player 1's payoffs satisfy increasing differences, so player 1 should use a cutoff strategy, doing \hat{T} if $t_1 > \theta_1$, doing B if $t_1 < \theta_1$, where θ_1 is some number between 0 and 1.

Then player 2 would think that the probability of 1 doing \hat{T} is $\text{Prob}(t_1 > \theta) = 1 - \theta$.

You can easily verify that there is no equilibrium where player 2 is sure to choose either L or R.

For player 2 to be willing to randomize between L and R, both L and R must give her the same expected payoff, so $0 = (-1)(1 - \theta_1) + (3)\theta_1$, and so $\theta_1 = 0.25$.

So in equilibrium, player 1 must use the strategy: do \hat{T} if $t_1 > 0.25$, do B if $t_1 < 0.25$.

For player 1 to be willing to implement this strategy, he must be indifferent between T and B when his type is exactly $t_1 = \theta_1 = 0.25$. Let q denote the probability of player 2 doing L.

Then to make type θ_1 indifferent between \hat{T} and B, q must satisfy $\varepsilon\theta_1 = (1)q + (-1)(1 - q)$, which implies $q = (1 + \varepsilon\theta_1)/2 = (1 + 0.25\varepsilon)/2$. (So as $\varepsilon \rightarrow 0$, q approaches 0.5.)

Now consider a game with two-sided incomplete information.

Suppose player 1's type t_1 is drawn from a Uniform distribution on the interval from 0 to 1, player 2's type t_2 is drawn independently from a Uniform distribution on the interval from 0 to 1, and the payoffs depend on 1's type as follows, for some given number ε between 0 and 1:

	L	R
\hat{T}	$\varepsilon t_1, \varepsilon t_2$	$\varepsilon t_1, -1$
B	$1, \varepsilon t_2$	$-1, 3$

With increasing differences, the action \hat{T} becomes more attractive to higher types of player 1.

Similarly, the action L becomes more attractive to higher types of player 2.

So we should look for an equilibrium where each uses a cutoff strategy of the form

- player 1 does \hat{T} if $t_1 > \theta_1$, player 1 does B if $t_1 < \theta_1$,
- player 2 does L if $t_2 > \theta_2$, player 2 does R if $t_2 < \theta_2$,

for some pair of cutoffs θ_1 and θ_2 .

It is easy to check that neither player's action can be certain to the other, and so these cutoffs θ_1 and θ_2 must be strictly between 0 and 1.

With t_1 Uniform on 0 to 1, the probability of player 1 doing \hat{T} ($t_1 > \theta_1$) is $1 - \theta_1$.

Similarly, the probability of player 2 doing L ($t_2 > \theta_2$) is $1 - \theta_2$.

The cutoff types must be indifferent between the two actions. So we have the equations

$$\varepsilon\theta_1 = (1)(1 - \theta_2) + (-1)\theta_2, \quad \varepsilon\theta_2 = (-1)(1 - \theta_1) + (3)\theta_1.$$

The unique solution to these equations is $\theta_1 = (2 + \varepsilon)/(8 + \varepsilon^2)$, $\theta_2 = (4 - \varepsilon)/(8 + \varepsilon^2)$.

Unless a player's type exactly equals the cutoff (which has zero probability), he is not indifferent between his two actions, and he uses the action yielding a higher expected payoff given his type.

As $\varepsilon \rightarrow 0$, these equilibria approach the randomized strategies $(.75[\hat{T}] + .25[B], .5[L] + .5[R])$.

These examples show how randomized equilibria can become pure-strategy equilibria in Bayesian games where each player has minor private information that determines his optimal action in equilibrium.

This is called purification of randomized equilibria by Bayesian games (John Harsanyi, *IJGT*, 1973.)

Action-probabilities and Belief Probabilities (from Osborne's Section 10.4)

Suppose that we are given some extensive game with imperfect information.

With imperfect information, each decision node now needs two labels, one indicating which player makes the decision here, and another indicating what is the state of the player's information at this node. In a game tree, I indicate these by a label of the form "Player.InformationState", where a player cannot distinguish his nodes that have the same information state. Nodes with the same player and information state form an information set, which is also commonly indicated by a dashed line or curve in the game tree. Two nodes with the same player and information label must be followed by the same set of feasible moves.

Given any randomized strategy for any player i , at any information set of player i that could occur with positive probability when he plays this strategy, we can compute a probability distribution over the set of possible actions for player i at this information set.

These probabilities are called action probabilities or move probabilities.

That is, the action-probability for any action c at any information state s of any player i , which we may denote by $\sigma_i(c|s)$, denotes the conditional probability that player i will choose action c if information set s occurs in the game. (I usually put action probabilities in parentheses "(•)".)

A behavioral strategy for player i is a list of an action-probability distribution for each of player i 's information sets.

A behavioral-strategy profile σ is a list of a behavioral strategy for each player, specifying an action probability $\sigma_i(c|s)$ for every possible action c at every possible information set s of every player i .

Given a behavioral-strategy profile σ , the prior probability $P(x|\sigma)$ of any node x in the game tree is the multiplicative product of all chance-probabilities and action-probabilities on the path that leads to this node from the starting node. (The chance probabilities on all branches that follow chance nodes are part of the given structure of the extensive game, from Osborne's Section 7.6.)

For any node x in any information set s of any player i , i 's belief-probability of x at s , which we may denote by $\mu_i(x|s)$, is the conditional probability of node x being true that player i would believe if some node in the information set s occurred in the game. (I put belief probabilities in angle brackets "<•>".)

By Bayes's formula, when player i moves at his information set s , the belief probability that player i should assign to any node x in this information set s should be:

$$\mu_i(x|s) = P(x|\sigma) / \sum_{y \in s} P(y|\sigma)$$

whenever this formula is well-defined (not 0/0). In the denominator of this formula, the sum is over all nodes y in the information set s , that is, over all nodes y that share the same "i.s" label as node x .

These probabilities would all be strictly positive if the behavioral strategy profile σ had full support, that is, if we had $\sigma_i(c|s) > 0$ for every possible action c at every information state s of every player i .

A belief system μ is a list of such belief probability distributions over the nodes of each information set of each player in the game.

A belief system μ is consistent (in the weak sense) with a behavioral-strategy profile σ if the beliefs satisfy all Bayes's formula, as above, whenever this formula is well-defined (not 0/0).

So weak consistency, as defined here, does not restrict beliefs at information sets that have zero probability.

A belief system μ is fully consistent with a behavioral strategy profile σ if μ is the limit of beliefs that would be consistent with a sequence of full-support strategy profiles that converge to σ .

A behavioral-strategy profile σ is sequentially rational given a belief system μ if, at every information set, the player is assigning positive probability only to actions that maximize his expected payoff, given his beliefs about the current node in his information set and given what the behavioral-strategy profile specifies about players' behavior after this information set.

A sequential equilibrium is a behavioral-strategy profile and a belief system such that the strategy profile is sequentially rational given the belief system, and the belief system is consistent with the strategy profile.

Sequential rationality may determine beliefs in zero-probability events, even if weak consistency does not!

Trading between a buyer and a seller, who knows more about the object being sold

Facts about Uniform distributions. Suppose that \tilde{X} is a random variable drawn from a Uniform distribution on the interval from A to B, for some given numbers A and B such that $A < B$.

Then $E(\tilde{X}) = (A+B)/2$. Furthermore, for any number θ between A and B:

$$\Pr(\tilde{X} < \theta) = \Pr(\tilde{X} \leq \theta) = (\theta - A)/(B - A),$$

$$E(\tilde{X} | \tilde{X} \leq \theta) = E(\tilde{X} | \tilde{X} < \theta) = (A + \theta)/2,$$

$$E(\tilde{X} | \tilde{X} \geq \theta) = E(\tilde{X} | \tilde{X} > \theta) = (\theta + B)/2.$$

Example. To illustrate the problems of trading between individuals who have different information, consider the following simple situation, involving two players.

Player 1 is the seller of some unique object which he owns.

Player 2 is the only possible buyer of this object.

Depending on the object's quality, it may be worth as little as \$40 to player 1 and \$60 to player 2 (if its quality is low) or as much as \$100 to player 1 and \$120 to player 2 (if its quality is high).

Player 1 knows the quality of the object. Let 1's type \tilde{t}_1 denote his value of keeping the object.

With any quality, the object would be worth \$20 more to player 2 than to player 1.

That is, given 1's type \tilde{t}_1 , the value of the object to player 2 would be $V_2(\tilde{t}_1) = \tilde{t}_1 + 20$.

Player 2's belief about \tilde{t}_1 is described by a Uniform distribution on the interval \$40 to \$100.

(So $E(\tilde{t}_1) = (40+100)/2 = 70$ and $E(V_2(\tilde{t}_1)) = E(\tilde{t}_1 + 20) = 90$.)

Game where buyer bids Suppose first that player 2 can offer to buy for any positive price r , and then player 1 will accept or reject the offer. If the offer is rejected then they each get profit 0.

If the offer is accepted then 1's profit is $r - \tilde{t}_1$ and 2's profit is $V_2(\tilde{t}_1) - r$.

In a subgame-perfect equilibrium, player 1 will accept if $\tilde{t}_1 < r$, but player 1 will reject if $\tilde{t}_1 > r$.

Player 2's expected profit from offering any price r is $\Pr(\tilde{t}_1 < r) (E(V_2(\tilde{t}_1) | \tilde{t}_1 < r) - r)$.

For any number r between 40 and 100, this expected profit is

$$\Pr(\tilde{t}_1 < r) (E(\tilde{t}_1 + 20 | \tilde{t}_1 < r) - r) = \Pr(\tilde{t}_1 < r) (E(\tilde{t}_1 | \tilde{t}_1 < r) + 20 - r) =$$

$$((r-40)/(100-40))((40+r)/2 + 20 - r) = (r-40)(80-r)/120 = (-3200 + 120r - r^2)/120.$$

This quadratic formula is maximized by letting $r = 60$.

(The buyer cannot gain by bidding less than 40 or more than 100, because a bid below 40 would be surely rejected, and a bid above 100 would be worse than the surely-accepted bid of 100.)

So in the unique subgame-perfect equilibrium of this game, player 2 offers to buy for \$60, and

player 1 accepts if $\tilde{t}_1 < 60$. The probability of trade is $\Pr(\text{trade}) = (60-40)/(100-40) = 1/3$.

(We could model 2 bidding as a game with perfect information where 1 learns \tilde{t}_1 after 2 chooses r .)

Game where seller bids Suppose now that player 1 can offer to buy for any positive price r , and then player 2 will accept or reject the offer. If the offer is rejected then they each get profit 0.

If the offer is accepted, then 1's profit is $r - \tilde{t}_1$ and 2's profit is $V_2(\tilde{t}_1) - r$.

In this game, the price is named by the player who has private information, and so signaling effects will give us many equilibria.

Let's look first for an equilibrium where there is some price \bar{r} such that player 2 would surely accept an offer to sell for \bar{r} but would surely reject an offer to sell for any price higher than \bar{r} .

In this equilibrium, player 1 will offer \underline{r} if $\tilde{t}_1 < \bar{r}$.

For player 2 to accept the offer \bar{r} , 2's expected profit from accepting \bar{r} must not be negative, so $0 \leq E(V_2(\tilde{t}_1) | \tilde{t}_1 < \bar{r}) - \bar{r} = E(\tilde{t}_1 + 20 | \tilde{t}_1 < \bar{r}) - \bar{r} = (40 + \bar{r})/2 + 20 - \bar{r}$, which implies $\bar{r} \leq 80$.

Player 2 can be expected to reject any offer to sell at a price greater than \bar{r} , because such a trade would be unprofitable for player 2 if she made the worst inference about player 1, which is that his type is $\tilde{t}_1 = 40$, in which case the object is worth only $40 + 20 = \$60$ to player 2.

So we can construct such an equilibrium for any \bar{r} such that $60 \leq \bar{r} \leq 80$.

In such an equilibrium, types higher than \bar{r} may be expected to make some offer higher than 120, which player 2 could never profitably accept.

An offer between \bar{r} and 120 may be rejected by player 2 because this surprise offer may lead player 2 to believe that 1's type is 40, in which case the object is only worth 60 to player 2.

Among these almost-pooling equilibria, player 1 most prefers the equilibrium with $\bar{r} = 80$.

In this equilibrium, the probability of trade is $\Pr(\text{trade}) = \Pr(\tilde{t}_1 < 80) = (80 - 40)/(100 - 40) = 2/3$.

There are many other equilibria where 1's types make more offers.

Let's look for an equilibrium in which some types of player 1 would offer to sell for \$70, but all higher types would offer to sell for \$100, and player 2 would be sure to accept \$70 but her probability of accepting \$100 would be between 0 and 1.

To find this equilibrium, there are unknowns that we must find:

let q denote the probability that player 2 would accept an offer of \$100,

and let θ denote the highest type of player 1 that would offer \$70.

For player 2 to be willing to randomize between accepting and rejecting \$100, her expected profit from accepting it must be 0, and so

$$0 = E(V_2(\tilde{t}_1) | \tilde{t}_1 > \theta) - 100 = E(\tilde{t}_1 + 20 | \tilde{t}_1 > \theta) - 100 = (\theta + 100)/2 + 20 - 100, \text{ and so } \theta = 60.$$

For player 1 to offer \$70 below when his type is below θ but \$100 when his type is above θ , we need that $70 - t_1 \geq q(100 - t_1)$ when $t_1 < \theta$, and $70 - t_1 \leq q(100 - t_1)$ when $t_1 > \theta$.

These inequalities imply $70 - \theta = q(100 - \theta)$, and so $q = (70 - 60)/(100 - 60) = 1/4$.

In this equilibrium, $\Pr(\text{trade}) = \Pr(\tilde{t}_1 < \theta) + \Pr(\tilde{t}_1 > \theta)q = (20/60) + (40/60)(1/4) = 1/2$.

(Advanced result: There is a separating equilibrium in which each possible type t_1 of player 1 would offer to sell for $r(t_1) = t_1 + 20$, and the probability q of player 2 accepting would depend on the offer r according to the formula $q(r) = e^{-(r-60)/20}$, for any $r \geq 60$.)

The Holdup Problem Player 1 can invest to improve an asset which he may later sell player 2.

To give the asset any quality $x \geq 0$, player 1 would have to make an investment that would cost him $c(x) = x^2$.

With this quality, the asset will be worth $v_1(x) = x$ to player 1, but it will be worth $v_2(x) = 2x$ to player 2.

We consider two different versions of this game, which differ in how they bargain over the price.

Buyer-offer game First player 1 chooses the quality $x \geq 0$. Player 2 observes this quality x .

Then player 2 chooses a price $p \geq 0$ at which she offers to buy the asset from player 1.

Player 1 observes this offer, and then can choose to accept or reject it. Final payoffs are:

$$u_1(x, p, \text{accept}) = p - c(x) = p - x^2, \quad u_2(x, p, \text{accept}) = v_2(x) - p = 2x - p,$$

$$u_1(x, p, \text{reject}) = v_1(x) - c(x) = x - x^2, \quad u_2(x, p, \text{reject}) = 0.$$

There is a unique subgame-perfect equilibrium.

At the last stage, given p and x , player 1 accepts if $p > v_1(x)$ and rejects if $p < v_1(x)$.

So player 2's optimal offer, given x , must be to offer $p = v_1(x) = x$, which player 1 must accept.

Note: In the case where $p = v_1(x)$, player 1 is actually indifferent between accepting and rejecting.

But if player 1 had any chance of rejecting in this case of indifference, then player 2 would instead want to offer the smallest p satisfying $p > v_1(x)$, to get 1's sure acceptance, and such a minimal p cannot be found! Thus, in a subgame-perfect equilibrium, player 1 must always accept when $p = v_1(x)$.

So player 1 knows that his payoff from quality x will be $v_1(x) - c(x) = x - x^2$, which is maximized by $x = 0.5$.

So the equilibrium outcome is: player 1 chooses quality $x = 0.5$, player 2 offers price $p = x = 0.5$,

and the players' payoffs are $u_1 = p - c(x) = 0.5 - (0.5)^2 = 0.25$, $u_2 = v_2(x) - p = 2 \times 0.5 - 0.5 = 1 - 0.5 = 0.5$.

Seller-offer game. First player 1 chooses the quality $x \geq 0$.

Then player 1 chooses the price $p \geq 0$ at which he offers to sell the asset.

Player 2 observes x and p , and then can choose to accept or reject 1's offer. Payoffs are still

$$u_1(x, p, \text{accept}) = p - x^2, \quad u_2(x, p, \text{accept}) = 2x - p, \quad u_1(x, p, \text{reject}) = x - x^2, \quad u_2(x, p, \text{reject}) = 0.$$

In the unique subgame-perfect equilibrium, player 2 accepts if $p \leq v_2(x)$ but rejects if $p > v_2(x)$,

so given any $x \geq 0$, player 1 offers $p = v_2(x) = 2x$. So player 1 chooses $x = 1$ to maximize $2x - x^2$.

Again, in the case when $p = v_2(x)$, player 2 would actually be indifferent between accepting and rejecting; but there would be no optimal p for player 1 if player 2 had any chance of rejecting in this case of indifference, as 1 would then want to offer the largest p such that $p < v_2(x)$, and no such number exists!

Such situations often arise in extensive games with perfect information when a player i can choose a number a_i in some interval with a continuum of possible values, and then another player j will make a decision next after observing i 's choice of a_i . It may happen that player j 's optimal decision would jump discontinuously from one alternative to another when player i choice is some critical value \bar{a}_i , and player j would be indifferent between these two alternatives when i 's choice is exactly equal to this value \bar{a}_i . In such cases, the only way to construct a subgame-perfect equilibrium may be to stipulate that, among the alternatives that are optimal for j when i has chosen \bar{a}_i , the alternative that player j would actually choose after \bar{a}_i must be the alternative that player i would prefer. Otherwise, player i would want to deviate infinitesimally from \bar{a}_i in the direction that would induce this preferred response by player j .

So the equilibrium outcome is: 1 chooses quality $x = 1$ and offers price $p = 2x = 2$,

and the players' payoffs are $u_1 = 2 - (1)^2 = 1$, $u_2 = 2(1) - 2 = 0$.

Notice that the equilibrium sum of payoffs $u_1 + u_2$ is greater in the seller-offer game ($1 + 0 > 0.25 + 0.5$).

That is, for an efficient outcome, the person who made the first-period investment should have more control in the process of bargaining over the price. If they were about to play the buyer-offer game, the buyer would be willing to sell her right to set the price for any payment more than 0.5, and the seller would be willing to pay up to 0.75 for the right to set the price.

Both of these games have many other Nash equilibria that are not subgame-perfect. Consider any (\hat{x}, \hat{p}) such that $v_2(\hat{x}) \geq \hat{p} \geq c(\hat{x}) + \max_{x \geq 0} (v_1(x) - c(x)) = c(\hat{x}) + 0.25$ (such as $\hat{x} = 1$, $\hat{p} = 1.625$), so that each does better than he or she could do alone. With either player offering the price, there is a Nash equilibrium in which 1 chooses this quality \hat{x} , and then this price \hat{p} is offered and accepted, but rejection would follow any other quality $x \neq \hat{x}$ or any other price-offer $p \neq \hat{p}$. These Nash equilibria violate sequential rationality, however, as threats to reject prices between $v_1(x)$ and $v_2(x)$ would not be credible.