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 $u: C \times S \to \mathbb{R}$, 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 $u: C \times S \to \mathbb{R}$ 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 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 $5 < 0.7\times7+0.3\times2 = 5.5$ and $3 < 0.7\times2+0.3\times7 = 3.5$.

Linear Duality Theorem (Farkas's lemma, theorem of the alternatives)

Given any m×n matrix $A = (a_{ij})_{i\in\{1,\dots,m\},j\in\{1,\dots n\}}$ and any vector $b = (b_i)_{i\in\{1,\dots,m\}}$ in $\mathbb{R}^m$, exactly one of the following two conditions is true:

(1) $\exists x \in \mathbb{R}^n$ such that $Ax \geq b$.

(2) $\exists y \in \mathbb{R}^m$ such that $y \geq \mathbf{0}$, $y'A = \mathbf{0}$, and $y'b > 0$.

Here $\mathbf{0}$ denotes a vector of zeroes in some appropriate number of dimensions.

Vector inequalities denote systems of numerical inequalities:

$Ax \geq b$ means: $\sum_{j=1}^{n} a_{ij}\, x_j \geq b_i \;\; \forall i\in\{1,\dots,m\}$,

$y'A = \mathbf{0}$ means $\sum_{i=1}^{m} y_i\, a_{ij} = 0 \;\; \forall j\in\{1,\dots,n\}$,

$y'b > 0$ means $\sum_{i=1}^{m} y_i\, b_i > 0$,

$y \geq \mathbf{0}$ means $y_i \geq 0 \;\; \forall i\in\{1,\dots,m\}$.

We may let $\mathbb{R}^m_+ = \{y\in\mathbb{R}^m |\; y\geq\mathbf{0}\}$ denote the nonnegative orthant in $\mathbb{R}^m$.

$y'(Ax-b) = \sum_{i=1}^{m} y_i \left((\sum_{j=1}^{n} a_{ij}\, x_j) - b_i\right)$.

Proof. Conditions (1) and (2) cannot both be true for any x and y,
because $y \geq \mathbf{0}$ and $Ax \geq b$ would imply $y'(Ax-b) \geq 0$,
while $y'A = \mathbf{0}$ and $y'b > 0$ would imply $y'(Ax-b) < 0$, a contradiction.
So (2) must be false if (1) is true.

Now suppose that (1) is false. This hypothesis means that the vector b is not in the set $\{Ax-z |\; x \in \mathbb{R}^n,\; z \in \mathbb{R}^m,\; z \geq \mathbf{0}\}$.
This set is convex and closed. So by the separating hyperplane theorem (MWG p948), there must exist some $y\in\mathbb{R}^m$ such that
$y'b > \max\{y'(Ax-z) |\; x\in\mathbb{R}^n, z\in\mathbb{R}^m, z\geq\mathbf{0}\}$.
This max must be nonnegative (because x and z could be $\mathbf{0}$), and it must be finite.
[In fact, this max must be exactly 0, because if we could achieve any $0 < \alpha = y'(Ax-z)$ with $x\in\mathbb{R}^n$, $z\in\mathbb{R}^m$, $z\geq\mathbf{0}$, then doubling x and z could achieve 2α, and so the max would be $+\infty$.]
If the component y_i were negative, then choosing large positive z_i could drive this max to $+\infty$.
If the j'th component of $y'A$ were nonzero, then choosing x_j with the same sign and large absolute value could also drive this max to $+\infty$.
So we must have $y'b > 0$, $y'A = \mathbf{0}$, and $y \geq \mathbf{0}$.
So (2) must be true if (1) is false.

Thus, (2) is true if and only if (1) is false. QED.

First duality application: strong domination. Suppose that there is a finite set of choice alternatives C, a finite set of possible states S, and u(c,s) denotes the utility payoff for the decision-maker if the state is s. Consider any given choice alternative d in C. By duality, exactly one of the following two conditions is true. *The dual variable for each constraint is shown at right in red italics:*

(1) $\exists p \in \mathbb{R}^S$ such that $\sum_{s\in S} (u(d,s) - u(c,s))\, p(s) \geq 0 \;\; \forall c\in C,$ *σ(c)*

$p(s) \geq 0 \;\; \forall s\in S,$ *δ(s)*

and $\sum_{s\in S} p(s) \geq 1.$ *α*

(2) $\exists(\sigma,\delta,\alpha) \in \mathbb{R}_+^C \times \mathbb{R}_+^S \times \mathbb{R}_+$ such that $\sum_{c\in C} \sigma(c)(u(d,s) - u(c,s)) + \delta(s) + \alpha = 0 \;\; \forall s\in S,$ *p(s)*

and $\alpha > 0.$

Condition (1) holds iff there is some probability distribution on the states such that d maximizes expected utility (because we could divide p by its sum to make it a probability distribution). Condition (2) holds iff there exists some probability distribution σ on the set of choice alternatives (a randomized strategy) such that $u(d,s) < \sum_{c\in C} \sigma(c)\, u(c,s) \;\; \forall s\in S$. (When (2) holds, the σ(c) cannot all be zero, and so we could divide (σ,δ,α) by $\sum_c \sigma(c)$ to make σ a probability distribution.) So d is optimal for some beliefs iff it is not strongly dominated by some randomized strategy. (Myerson, 1991, Theorem 1.6)

Second application: Utility theory. Suppose there are n possible prizes numbered 1,...,n. We ask a decision maker m questions. In the i'th question, we ask whether he prefers a lottery p(i), which offers each prize j with probability $p_j(i)$, or a lottery q(i) which offers each prize j with probability $q_j(i)$. In each case, suppose p(i) denotes the lottery that he strictly prefers. Let ε be any positive number. By duality, exactly one of the following two conditions is true:

(1) $\exists u \in \mathbb{R}^n$ such that $\sum_{j=1}^n (p_j(i) - q_j(i))\, u_j \geq \varepsilon \;\; \forall i\in\{1,...,m\}.$ *σ_i*

(2) $\exists \sigma \in \mathbb{R}_+^m$ such that $\varepsilon \sum_{i=1}^m \sigma_i > 0,$ and $\sum_{i=1}^m \sigma_i\, (p_j(i) - q_j(i)) = 0 \; \forall j\in\{1,...,n\}.$ *u_j*

Condition (1) holds iff we can find a utility function for which his revealed preferences are compatible with expected utility maximization. Condition (2) holds iff we can find a violation of the substitution axiom in compound lotteries. (Renormalize so that $\sum_i \sigma_i = 1$, then consider the first compound lottery that gives each p(i) lottery with probability σ_i, and the second compound lottery that gives each q(i) lottery with probability σ_i. After the first stage of the compound lotteries, the first would always seem better than the second, but the ex ante probability of each prize j is equal in the two compound lotteries.) So the revealed preferences are compatible with expected utility maximization iff there is no violation of the substitution axiom.

Third application: weak domination. Let S, C, u, and d be as in the previous example. Let ε be any small positive number. By duality, exactly one of the following two conditions is true.

(1) $\exists p \in \mathbb{R}^S$ such that $\sum_{s\in S} (u(d,s) - u(c,s))\, p(s) \geq 0 \;\; \forall c\in C,$ *σ(c)*

and $p(s) \geq \varepsilon \;\; \forall s\in S.$ *δ(s)*

(2) $\exists(\sigma,\delta) \in \mathbb{R}_+^C \times \mathbb{R}_+^S$ such that $\sum_{c\in C} \sigma(c)(u(d,s) - u(c,s)) + \delta(s) = 0 \;\; \forall s\in S,$ *p(s)*

and $\varepsilon \sum_{s\in S} \delta(s) > 0.$

Condition (1) holds iff there is some probability distribution p on the set of states such that every state has strictly positive probability and d maximizes expected utility over all choices in C. Condition (2) holds iff there exists some probability distribution σ on the set of choice alternatives such that $u(d,s) \leq \sum_{c\in C} \sigma(c)\, u(c,s) \;\; \forall s\in S$, with at least one strict inequality (<) for some s. So d is optimal for some beliefs where all states have positive probability iff d is not weakly dominated by some randomized strategy. (Myerson, 1991, Theorem 1.7)

Separating Hyperplane Theorem (MWG M.G.2):
Suppose X is a closed convex subset of $\mathbb{R}^N$, and w is a vector in $\mathbb{R}^N$.
Then exactly one of the following two statements is true: Either (1) $w \in X$,
or (2) there exists a vector $y\in \mathbb{R}^N$ such that $y'w > \max_{x\in X} y'x$
(but not both). (Here $y'w = y_1x_1+...+y_Nx_N$, with $y = (y_1,...,y_N)$ and $x = (x_1,...,x_N)$.)

Supporting Hyperplane Theorem (MWG M.G.3):
Suppose X is a convex subset of $\mathbb{R}^N$, and w is a vector in $\mathbb{R}^N$. Then exactly one of the following two statements is true: Either (1) w is in the interior of X (relative to $\mathbb{R}^N$),
or (2) there exists a vector $y\in \mathbb{R}^N$ such that $y\neq\mathbf{0}$ and $y'w \geq \max_{x\in X} y'x$
(but not both). Here $\mathbf{0} = (0,...,0)$.

Fact If X is convex and $\max_{x\in X} y'x$ is a finite number, then this maximum must be achieved at some extreme point in X. (MWG p 946.)

Strong domination Theorem. Given the nonempty finite sets C={choices}, S={states}, the utility function $u:C\times S\to\mathbb{R}$, and the choice $d\in C$, exactly one of these two statements is true:
Either (1) $\exists\sigma\in\Delta(C)$ such that $u(d,s) < \sum_{c\in C} \sigma(c)u(c,s)$ $\forall s\in S$,
or (2) $\exists p\in\Delta(S)$ such that $\sum_{s\in S} p(s)u(d,s) = \max_{c\in C} \sum_{s\in S} p(s)u(c,s)$.

Proof. Let $X = \{x\in\mathbb{R}^S | \exists\sigma\in\Delta(C) \text{ s.t. } x_s \leq \sum_{c\in C} \sigma(c)u(c,s)\ \forall s\}$. X is a convex subset of $\mathbb{R}^S$.
Condition (1) here is equivalent to: (1′) the vector $u(d) = (u(d,s))_{s\in S}$ is in the interior of X.
By the Supporting Hyperplane Thm, (1′) is false iff
(2′) $\exists p\in\mathbb{R}^S$ such that $p\neq\mathbf{0}$ and $\sum_{s\in S} p(s)u(d,s) \geq \max_{x\in X} p(s)x_s$.
We must have $p(s)\geq 0$ for all s, because x in X can have x_s approaching $-\infty$.
So $p\geq\mathbf{0}$ and $p\neq\mathbf{0}$ implies $\sum_{s\in S} p(s)>0$. Dividing by this sum, we can make $\sum_{s\in S} p(s)=1$ (wlog).
Furthermore, the maximum of the linear function $p'x$ over $x\in X$ must be achieved at one of the extreme points in X, which are vectors $(u(c,s))_{s\in S}$ for the various $c\in C$.
So (2′) is equivalent to condition (2) in the theorem here.

Expected Utility Theorem. Let N be a finite set of prizes, and consider a finite sequence of pairs of lotteries $p(i)\in\Delta(N)$ and $q(i)\in\Delta(N)$, for $i\in M=\{1,...,m\}$. (Here $p(i) = (p_j(i))_{j\in N}$, and $q(i) = (q_j(i))_{j\in N}$.) Then exactly one of these two statements is true:
Either (1) $\exists\sigma\in\Delta(M)$ such that $\sum_{i\in M} \sigma(i)p_j(i) = \sum_{i\in M} \sigma(i)q_j(i)\ \forall j\in N$,
or (2) $\exists u\in\mathbb{R}^N$ such that $\sum_{j\in N} p_j(i)u_j > \sum_{j\in N} q_j(i)u_j\ \ \forall i\in M$.

Proof. Let $X = \{\sum_{i\in M} \sigma(i)(q(i)-p(i)) | \sigma\in\Delta(M)\}$. Then X is a closed convex subset of $\mathbb{R}^N$.
Condition (1) here is equivalent to: (1′) the N-vector $\mathbf{0}$ is in X.
By the Separating Hyperplane Thm, (1′) is false iff
(2′) $\exists u\in \mathbb{R}^N$ such that $0 = u'\mathbf{0} > \max_{x\in X} u'x$.
The extreme points of X are vectors $(q(i)-p(i)) = (q_j(i)-p_j(i))_{j\in N}$, and the linear function $u'x = \sum_{j\in N} x_ju_j$ must achieve its maximum over $x\in X$ at one of these extreme points.
So (2′) is equivalent to (2) in the theorem here.

Fact. Suppose the utility-representation condition (2) is satisfied by $u = (u_j)_{j\in N}$.
Then (2) is also satisfied by $\hat{u}$ iff there exists $A>0$ and B such that $\hat{u}_j = Au_j+ B\ \forall j\in N$.

A **strategic-form game** is characterized by $(N, (C_i)_{i\in N}, (u_i)_{i\in N})$ where
$N = \{1,2,...,n\}$ is the set of players, and, for each player i:
C_i is the set of alternative actions or (pure) strategies that are feasible for i in the game, and
$u_i:C_1\times C_2\times...\times C_n\rightarrow\mathbb{R}$ is player i's utility function in the game.
We generally assume that each player i independently chooses an action in C_i.
If $c = (c_1,c_2,...,c_n)$ is the combination (or profile) of actions chosen by the players
then each player i will get the expected utility payoff $u_i(c_1,c_2,...,c_n)$.
We let $C = C_1\times C_2\times...\times C_n = \times_{i\in N} C_i$ denote the set of all combinations or profiles of actions that the players could choose.
We let C_{-i} denote the set of all profiles of actions that can be chosen by players other than i.
When $c\in C$ is a profile of actions for the players, c_i denotes the action of each player i,
c_{-i} denotes the profile of actions for players other than i where they act as in c,
and $(c_{-i};d_i)$ denotes the profile of actions in which i's action is changed to d_i but all others choose the same action as in c. (We may use this notation even if player i is not the "last" player.)

A randomized strategy (or mixed strategy) for player i is a probability distribution over C_i,
so $\Delta(C_i)$ denotes the set of all randomized strategies for player i. (pure=nonrandomized.)
An action d_i for player i is strongly dominated by a randomized strategy $\sigma_i \in \Delta(C_i)$ if

$$u_i(c_{-i};d_i) < \sum_{c_i\in C_i} \sigma_i(c_i)u_i(c_{-i};c_i) \quad \forall c_{-i}\in C_{-i}.$$

An action d_i for player i is weakly dominated by a randomized strategy $\sigma_i \in \Delta(C_i)$ if

$$u_i(c_{-i};d_i) \leq \sum_{c_i\in C_i} \sigma_i(c_i)u_i(c_{-i};c_i) \quad \forall c_{-i}\in C_{-i},$$ with strict inequality (<) for at least one c_{-i}.

The set of player i's best responses to any profile of opponents' actions c_{-i} is

$$\beta_i(c_{-i}) = \text{argmax}_{d_i\in C_i} u_i(c_{-i};d_i) = \left\{d_i\in C_i \middle| u_i(c_{-i};d_i) = \max_{c_i\in C_i} u_i(c_{-i};c_i)\right\}.$$

Similarly, if i's beliefs about the other players' actions can be described by a probability distribution μ in $\Delta(C_{-i})$, then the set of player i's best responses to the beliefs μ is

$$\beta_i(\mu) = \text{argmax}_{d_i\in C_i} \sum_{d_{-i}\in C_{-i}} \mu(c_{-i})u_i(c_{-i};d_i).$$

Fact. If we iteratively eliminate strongly dominated actions for all players until no strongly dominated actions remain, then we get a reduced game in which each remaining (*rationalizable*) action for each player is a best response to some beliefs about the other players' actions.

If each player j independently uses a strategy σ_j in $\Delta(C_j)$, then player i's expected payoff is

$$u_i(\sigma_{-i};\sigma_i) = u_i(\sigma_1,\sigma_2,...,\sigma_n) = \sum_{c\in C} \left(\prod_{j\in N} \sigma_j(c_j)\right) u_i(c)$$
$$= \sum_{c_i\in C_i}\sigma_i(c_i)\sum_{c_{-i}\in C_{-i}}\left(\prod_{j\in N-i}\sigma_j(c_j)\right)u_i(c_{-i};c_i) = \sum_{c_i\in C_i}\sigma_i(c_i)u_i(\sigma_{-i};[c_i]).$$

Fact $\sigma_i \in \text{argmax}_{\tau_i\in\Delta(C_i)} u_i(\sigma_{-i};\tau_i)$ if and only if σ_i assigns positive probability $\sigma_i(c_i)>0$ only to actions c_i that are in $\text{argmax}_{d_i\in C_i} u_i(\sigma_{-i};[d_i])$.
A Nash equilibrium is a profile of actions or randomized strategies such that each player is using a best response to the others. That is $\sigma = (\sigma_1,...,\sigma_n)$ is a Nash equilibrium in randomized strategies iff $\sigma_i \in \text{argmax}_{\tau_i\in\Delta(C_i)} u_i(\sigma_{-i};\tau_i)$ for every player i in N.
Fact. Any finite strategic-form game has at least one Nash equilibrium in randomized strategies.

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. We describe here a procedure for finding Nash equilibria, from section 3.3 of Myerson (1991).

We are given some game, including a given set of players N and, for each i in N, a given set of feasible actions C_i for player i and a given payoff function $u_i:C_1\times...\times C_n\rightarrow\mathbb{R}$ for player i.
The <u>support</u> 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 5-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 C_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: $\sum_{si\in Si}\sigma_i(s_i) = 1\ \ \forall i\in N$; and $Eu_i(s_i|\sigma_{-i}) = w_i\ \ \forall i\in N\ \ \forall s_i\in S_i$.
(Here $\forall$ means "for all", $\in$ means "in".) We have as many equations as unknowns (w_i, $\sigma_i(s_i)$).

(3) 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 (4) and (5)

(4) 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.

(5) A solution from (2) that satisfies the condition in (4) 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 C_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.
Now, for every action a_i that is in C_i but is not in S_i (so $\sigma_i(a_i)=0$), 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.

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

Player 1	Player 2: L	M	R
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 $Eu_1(T|\sigma_2) = Eu_1(B|\sigma_2)$, and so
$7q+2(1-q-r)+3r = 2q+7(1-q-r)+4r$.
Player 2 randomizing over {L,M,R} requires $Eu_2(L|\sigma_1) = Eu_2(M|\sigma_1) = Eu_2(R|\sigma_1)$, and so
$2p+7(1-p) = 7p+2(1-p) = 6p+5(1-p)$.
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 $Eu_1(T|\sigma_2) = Eu_1(B|\sigma_2)$, so $2(1-r)+3r = 7(1-r)+4r$.
Player 2 randomizing over {M,R} requires $Eu_2(M|\sigma_1) = Eu_2(R|\sigma_1)$, so $7p+2(1-p) = 6p+5(1-p)$.
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 $Eu_1(T|\sigma_2) = Eu_1(B|\sigma_2)$, so $7q+2(1-q) = 2q+7(1-q)$.
Player 2 randomizing over {L,M} requires $Eu_2(L|\sigma_1) = Eu_2(M|\sigma_1)$), so $2p+7(1-p) = 7p+2(1-p)$.
These solution for these two equations in two unknowns is $p = 1/2$ and $q = 1/2$.
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 $Eu_2(R|\sigma_1) = 6p+5(1-p) = 6\times1/2+5\times1/2 = 5.5 > Eu_2(L|\sigma_1) = 2\times1/2+7\times1/2 = 3.5$.
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]. (q=1−r)
Player 1 randomizing over {T,B} requires $Eu_1(T|\sigma_2) = Eu_1(B|\sigma_2)$, so $7q+3(1-q) = 2q+4(1-q)$.
Player 2 randomizing over {L,R} requires $Eu_2(L|\sigma_1) = Eu_2(R|\sigma_1)$, so $2p+7(1-p) = 6p+5(1-p)$.
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;
$Eu_2(M|\sigma_1) = 7p+2(1-p) = 7\times1/3+2\times2/3 = 11/3 < Eu_2(L|\sigma_1) = 2\times1/3+7\times2/3 = 16/3$.
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 $Eu_1 = 7\times1/6+3\times5/6 = 2\times1/6+4\times5/6 = 11/3 = 3.667$
and $Eu_2 = 2\times1/3+7\times2/3 = 6\times1/3+5\times2/3 = 16/3 = 5.333$.

Increasing differences and increasing strategies in Bayesian games

We may consider Bayesian games where first each player i first learns his type $\tilde{t}_i$, and then each player i chooses his action a_i. We will assume that each player i's type is drawn from some probability distribution p_i, independently of all other players' types, so that $p_i(t_i) = \text{Prob}(\tilde{t}_i = t_i)$. The payoffs of each player i may depend on all players' types and actions according to some utility payoff function $U_i(c_1,...,c_n,\tilde{t}_1,...,\tilde{t}_n)$.

Consider a two-player Bayesian game where player 1 has two possible actions, 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 when he switches from B to T is
$U_1(T,c_2,t_1,t_2) - U_1(B,c_2,t_1,t_2)$
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 increasing differences if this difference
$U_1(T, c_2, t_1, t_2) - U_1(B, c_2, t_1, t_2)$ would increase (weakly, increase or stay constant) whenever player 1's type t_1 is increased, no matter what player 2's action c_2 and type t_2 may be.
That is, increasing differences means that, for every r_1, t_1, c_2, and t_2:
if $r_1 \geq t_1$ then $U_1(T, c_2, r_1, t_2) - U_1(B, c_2, r_1, t_2) \geq U_1(T, c_2, t_1, t_2) - U_1(B, c_2, t_1, t_2)$.
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.
So in a cutoff strategy, there is at most one type which is randomizing between T and B, and all higher types must be choosing T for sure, and all lower types must be choosing B for sure.

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.
Notice that the probability of player 1 choosing T decreases as the cutoff θ increases.

More generally, in games where player 1's action can be any number in some range, we say that player 1's payoffs satisfy increasing differences if, for every pair of possible actions c_1 and d_1 such that $c_1>d_1$, the difference $U_1(c_1, c_2, t_1, t_2) - U_1(d_1, c_2, t_1, t_2)$ would increase whenever player 1's type t_1 is increased, no matter what player 2's action c_2 and type t_2 may be.
Fact. If 1's payoffs satisfy increasing differences then, for any strategy of player 2, player 1 will have a best-response strategy that is increasing, in the sense that 1's action would increase (weakly) as his type increases.
("Weakly" increasing here means that two different types might use the same action, or the higher type might use a higher action, but the higher type would never use a lower action.)

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 {T,B}, 2's actions are {L,R}.
Given 1's type t_1, the payoff matrix is

	L	R
T	t_1, 0	t_1, −1
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:
$U_1(T,L,t_1)-U_1(B,L,t_1) = t_1-1$, $U_1(T,R,t_1)-U_1(B,R,t_1) = t_1+1$.
Notice that these differences increase in t_1. So higher types t_1 always find T relatively more attractive than lower types, and player 1 will use a cutoff strategy.
The possible cutoff strategies are:

- (θ>.3) every type would choose [B], so 2 thinks the probability of T is P(T)=0;
- (θ=.3) {0,.1,.2} would choose [B], but .3 would randomize in some way, so P(T) is between 0 and 1/4;
- (.2<θ<.3) {0,.1,.2} would choose [B], but .3 would choose [T], so P(T) = 1/4;
- (θ=.2) {0,.1} would choose [B], .2 would randomize in some way, .3 would choose [T], and so P(T) is between 1/4 and 1/2;
- (.1<θ<.2) {0,.1} would choose [B], {.2,.3} would choose [T], so P(T) = 1/2;
- (θ=.1) 0 would choose [B], .1 would randomize in some way, {.2,.3} would choose [T], and so P(T) is between 1/2 and 3/4;
- (0<θ<.1) 0 would choose [B], {.1,.2,.3} would choose [T], so P(T) = 3/4;
- (θ=0) 0 would randomize in some way, {.1, .2,.3} would choose [T], and so P(T) is between 3/4 and 1;
- (θ<0) every type would choose [T], and so P(T) = 1.

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,
$P(T)(0) + (1-P(T))(0) = P(T)(-1) + (1-P(T))(3)$, and so P(T) = 3/4.
So the cutoff θ must be between 0 and .1 (0 would choose [B], {.1,.2,.3} would choose [T]).
Let q denote the probability that 2 chooses 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 1 prefer T when t_1 = .1.
$EU_1(T|t_1=0) \le EU_1(B|t_1=0)$ implies $(q)(0)+(1-q)(0) \le (q)(1)+(1-q)(-1)$, and so $1/2 \le q$.
$EU_1(T|t_1=.1) \ge EU_1(B|t_1=.1)$ implies $(q)(.1)+(1-q)(.1) \ge (q)(1)+(1-q)(-1)$, and so $q \le 11/20$.
So in equilibrium, 1 chooses B if t_1=0, 1 chooses T if $t_1 \ge .1$, and 2 randomizes, choosing L with some probability q that is between 1/2 and 11/20.

Now suppose instead player 1has five possible types {0, .1, .2, .3, .4}, each with probability 1/5.
To make player 2 willing to randomize, player 1 must use a strategy such that P(T) = 3/4.
For that to occur in an increasing cutoff strategy, the cutoff must be at θ=.1. So t_1=0 chooses B; and when t_1>.1 (which has probability 3/5) player 1 chooses T. The remaining 3/4−3/5 = 0.15 probability of T must come from 1 choosing T with probability (0.15/0.2) = 0.75 when t_1=.1 .
To make type t_1=.1 willing to randomize, player 2's probability of choosing L must be q=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 number between 0 and 1 (say ε=0.1):

	L	R
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 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 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 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 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 $\varepsilon \to 0$, q approaches 0.5.)

Now consider a game with two-sided incomplete information from Myerson (1991) section 3.10.
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
T	$\varepsilon t_1, \varepsilon t_2$	$\varepsilon t_1, -1$
B	$1, \varepsilon t_2$	$-1, 3$

With increasing differences, the action 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 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 us 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 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
$\varepsilon\theta_1 = (1)(1-\theta_2) + (-1)\theta_2$, $\varepsilon\theta_2 = (-1)(1-\theta_1) + (3)\theta_1$.
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 $\varepsilon \to 0$, these equilibria approach the randomized strategies (.75[T]+.25[B], .5[L]+.5[R]).

Introduction to auctions Consider n=2 bidders in a <u>first-price auction</u> to buy an object for which they have <u>independent private values</u> drawn from a Uniform distribution on 0 to M.
The set of players is N={1,2}. Each player i's type set is $T_i = [0,M]$, where the type t_i is i's value of the object being sold. Player i's action is a bid c_i which must be a nonegative number in $\mathbb{R}_+$.
The high bidder gets the object, which is worth his type to him, but the winner must pay the amount that he bid. Losers pay nothing. So the utility function for each player i is
$u_i(c_1,c_2,t_1,t_2) = t_i-c_i$ if $c_i>c_{-i}$, $u_i(c_1,c_2,t_1,t_2) = 0$ if $c_i<c_{-i}$, $u_i(c_1,c_2,t_1,t_2) = (t_i-c_i)/2$ if $c_i=c_{-i}$.
A strategy for player i specifies i's bid as some function of i's type, say $c_i = b_i(t_i)$.
Let us try to find a <u>symmetric</u> equilibrium of this game, and let us guess that the equilibrium strategy is <u>linear</u>, of the form $b_i(t_i) = \alpha t_i$ for some $\alpha>0$. Can this be an equilibrium, for some α?
Consider the problem of player 1's best response, when player 2 uses such a strategy, so $\tilde{c}_2 = \alpha\tilde{t}_2$.
When player 1 knows his type is t_1, player i's expected payoff from choosing bid c_1 would be
$EU_1(c_1|t_1) = (t_1-c_1)\,P(\tilde{c}_2 < c_1) = (t_1-c_1)\,P(\alpha\tilde{t}_2 < c_1) = (t_1-c_1)\,P(\tilde{t}_2 < c_1/\alpha) = (t_1-c_1)(c_1/\alpha)/M$,
(assuming that c_1 is between 0 and αM). So the first-order optimality conditions are
$0 = \partial EU_1(c_1|t_1)/\partial c_1 = (t_1-2c_1)/(\alpha M)$, which implies $c_1 = t_1/2$.
So player 1's best-response strategy is the same linear function as 2's strategy if $\alpha = 1/2$.
Thus, bidders 1 and 2 each bidding half of his or her type-value is an equilibrium in this auction.
E(payment from $i|\tilde{t}_i=t_i) = (t_i/2)\,P(\tilde{t}_{-i}/2 < t_i/2) = (t_i/2)(t_i/M) = t_i^2/(2M)$.

Now let us change the game to an <u>all-pay-own-bid auction</u>, with the same bidders and types.
As before, there are two bidders with independent private values drawn from Uniform [0,M], and the high bidder gets the object. But now each pays his own bid whether he wins or loses.
So now i's payoff is: $u_i = t_i-c_i$ if $c_i>c_{-i}$, $u_i = -c_i$ if $c_i<c_{-i}$, $u_i = t_i/2-c_i$ if $c_i=c_{-i}$.
There is no linear symmetric equilibrium to this game. (If 2 used the linear strategy $\tilde{c}_2 = \alpha\tilde{t}_2$ then 1's best response would be to bid $c_1=\alpha M$ if $t_1>\alpha M$ and $c_1=0$ if $t_1<\alpha M$, which is not linear!)
As our next guess, let us try to find a symmetric quadratic equilibrium where each $c_i = \alpha\, t_i^2$ for some $\alpha>0$. So to find this equilibrium, suppose that player 2 is using $\tilde{c}_2 = \alpha\,\tilde{t}_2^2$.
When player 1 knows his type is t_1, player i's expected payoff from choosing bid c_1 would be
$EU_1(c_1|t_1) = t_1\,P(\tilde{c}_2 < c_1) - c_1 = t_1\,P(\alpha\,\tilde{t}_2^2 < c_1) - c_1 = t_1\,P(\tilde{t}_2 < (c_1/\alpha)^{0.5}) - c_1$
$= t_1(c_1/\alpha)^{0.5}/M - c_1$ (assuming that c_1 is between 0 and αM^2).
So the first-order optimality conditions are $0 = \partial EU_1(c_1|t_1)/\partial c_1 = t_1/(2\alpha^{0.5}Mc_1^{0.5}) - 1$,
which implies $c_1 = t_1^2/(4\alpha M^2)$. For player 1's best-response strategy to be the same as the given strategy for player 2, we need $1/(4\alpha M^2) = \alpha$, which holds when $\alpha = 1/(2M)$.
So there is an equilibrium of this all-pay-own-price auction where each i bids $c_i = t_i^2/(2M)$.

Finally, let us change the game to a <u>second-price</u> auction. As in the first-price auction, the high bidder wins the object and is the only bidder to pay anything (the loser pays nothing), but now the amount that the high bidder pays is the second-highest bid, submitted by the other bidder.
So now i's payoff is: $u_i = t_i-c_{-i}$ if $c_i>c_{-i}$, $u_i = 0$ if $c_i<c_{-i}$, $u_i = (t_i-c_{-i})/2$ if $c_i=c_{-i}$.
In this auction, there is an equilibrium in which each bidder honestly bids his value $c_i = b_i(t_i) = t_i$.
In fact bidding $c_i=t_i$ weakly dominates any other strategy. Higher bids only add unprofitable wins where $c_{-i}>t_i$. Lower bids don't reduce what i pays and may lose profit opportunities when $c_{-i}<t_i$.
E(payment from $i|\tilde{t}_i=t_i) = E(\tilde{t}_{-i}|\tilde{t}_{-i}<t_i)\,P(\tilde{t}_{-i}<t_i) = (t_i/2)(t_i/M) = t_i^2/(2M)$

<u>Fact.</u> Given any type t_i, i's expected payment is $t_i^2/(2M)$ in these equilibria of different auctions.
Given i's type t_i, i's expected profit in each eqm is $EU_i(t_i) = t_i\,P(\tilde{t}_{-i}<t_i) - t_i^2/(2M) = t_i^2/(2M)$.
The seller's expected revenue from the bidders is $2E(\tilde{t}_i^2/(2M)) = 2\int_0^M (t^2/(2M))\,dt/M = M/3$.

A first-price auction with common values (from Myerson 1991, section 3.11).
Consider n=2 bidders in a first-price auction to buy an object which would have the same common value $\tilde{V}$ to either bidder, if he were to win it, but suppose that each bidder has different private information about this common value. Each bidder i observes a signal $\tilde{t}_i$ that is drawn independently from a uniform distribution on the interval from 0 to 1, and the value of the object to either bidder depends linearly on the signals according to the formula $\tilde{V} = A_1\tilde{t}_1+A_2\tilde{t}_2$.
So each player i's payoff $u_i(c_1,c_2,t_1,t_2)$ depends on the bids (c_1,c_2) and the types (t_1,t_2) as follows:
$u_i = A_1t_1+A_2t_2-c_i$ if $c_i>c_{-i}$, $u_i = 0$ if $c_i<c_{-i}$, $u_i = (A_1t_1+A_2t_2-c_i)/2$ if $c_i=c_{-i}$.
A strategy for player i specifies i's bid as some function of i's type, say $c_i = b_i(t_i)$.

Let us guess that the equilibrium strategies are linear, of the form $b_i(t_i) = \alpha_i t_i$ for some $\alpha_i>0$.
Nobody would want to bid more than the highest possible bid of his opponent, and so we must have $\alpha_1 = \alpha_2 = \alpha$, so that both bidders have the same range $[0,\alpha]$ of possible bids in equilibrium.
Now suppose that bidder 1, believing that 2 will bid $b_2(\tilde{t}_2) = \alpha\tilde{t}_2$, knows his type t_1 and is thinking of bidding some other c. Then player 1 will win if $\alpha\tilde{t}_2 < c$, that is, $\tilde{t}_2 < c/\alpha$.
Then 1's expected payoff would be
$$EU_1(c|t_1) = \int_0^{c/\alpha}(A_1t_1+A_2t_2-c)dt_2 = [A_1t_1 + 0.5A_2c/\alpha - c](c/\alpha)$$
$$= [E(\tilde{V}|\ \tilde{t}_1 = t_1, \alpha\tilde{t}_2 < c) - c]\ P(\alpha\tilde{t}_2 < c).$$
First-order conditions for c to be an optimal bid are then
$0 = \partial EU_1(c|t_1)/\partial c = [A_1t_1 + A_2c/\alpha - 2c]/\alpha$, and so $c = A_1t_1/(2-A_2/\alpha)$.
Thus, for 1's optimal bid here to be $c = \alpha t_1$, we need $\alpha = A_1/(2-A_2/\alpha)$,
and so our equilibrium must have $\alpha = 0.5(A_1+A_2)$.
(This symmetric formula also works for 2, who wants to bid αt_2 when 1 is expected to bid $\alpha\tilde{t}_1$.)
So the expected profit for type t_1 of player 1 when he bids $b_1(t_1) = \alpha t_1 = 0.5(A1+A_2)t_1$ in this equilibrium is $EU_1(t_1) = [A_1t_1 + 0.5A_2t_1 - 0.5(A_1+A_2)t_1]t_1 = 0.5A_1t_1^2$.

Let $\tilde{v}_i = A_i\tilde{t}_i$ denote the value that player i has privately seen going into the object here.
So in this model, each $\tilde{v}_i$ is an independent uniform random variable on the interval from 0 to A_i.
In terms of his privately observed value v_i, player i's equilibrium bid is $0.5(1+A_{-i}/A_i)v_i$,
and player i's conditional expected profit given his type is $0.5v_i^2/A_i$.

Example: Suppose $A_1=A_2=100$. If $\tilde{t}_1=0.01$ then 1 bids $100t_1=1$ and gets P(win)=0.01.
$E(\tilde{V}|\tilde{t}_1=0.01) = 51$, but $EU_1(c_1|\tilde{t}_1=0.01) = (1+0.5c_1 - c_1)(c_1/100) < 0$ if $c_1 > 2$.

Now let $\tilde{t}_0$ be another uniform [0,1] random variable that is observed by both bidders.
If we increased the common value by the commonly known amount $A_0\tilde{t}_0$, then the equilibrium bid for each type of each bidder would increase by this commonly known amount $A_0\tilde{t}_0$.
That is, if the common value were $\tilde{V} = A_0\tilde{t}_0+A_1\tilde{t}_1+A_2\tilde{t}_2$, where bidder 1 observes $\tilde{t}_0$ and $\tilde{t}_1$ and bidder 2 observes $\tilde{t}_0$ and $\tilde{t}_2$, then the equilibrium bidding strategies would be
$b_1(\tilde{t}_0,\tilde{t}_1) = A_0\tilde{t}_0 + 0.5(A_1+A_2)\tilde{t}_1$ and $b_2(\tilde{t}_0,\tilde{t}_2) = A_0\tilde{t}_0 + 0.5(A_1+A_2)\tilde{t}_2$.
Example: $A_0 = 1 = A_1$, $A_2 = 99$. Then equilibrium strategies are $b_i(\tilde{t}_0,\tilde{t}_1) = \tilde{t}_0 + 50\tilde{t}_i$, for i=1,2.
Notice that 1's two signals are equally minor in their impact on the value of the object,
but 1's bid depends much more on his private information $\tilde{t}_1$ than his shared information $\tilde{t}_0$.

Move probabilites, belief probabilities and sequential equilibria

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

Given any randomized strategy for any player i, at any information set t_i 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 $\{d_i\}$ for player i at this information set.

These probabilities $\sigma_i(d_i|t_i)$ are called <u>move probabilities</u> (or <u>action probabilities</u>).

That is, the move-probability for any move d_i at any information state t_i of any player i denotes the probability that player i will choose move d_i if information set t_i occurs in the game.

A <u>behavioral strategy</u> σ_i for player i is a vector that specifies a move-probability distribution for each of player i's information sets.

A <u>behavioral-strategy profile</u> σ is a vector that specifies a behavioral strategy σ_i for each player i, and so it must specify an move probability $\sigma_i(d_i|t_i)$ for every possible move d_i at every possible information set t_i of every player i in the game.

Given a profile of behavioral or randomized strategies for all players in the game, the <u>prior probability</u> P(x) of any node x in the tree is the multiplicative product of all chance-probabilities and move-probabilities on the path that leads to this node from the starting node.

(Here the chance probabilities on all branches that follow chance nodes are part of the given structure of the extensive game. We assume that these chance probabilities are all positive.)

A <u>full-support</u> behavioral strategy profile is one that assigns strictly positive probability ($\sigma_i(d_i|t_i)>0$) to every possible move d_i at every information set t_i of every player i, so that every node x in the game tree will have positive probability.

When player i moves at his information set t_i, the <u>belief probability</u> that player i should assign to any node x in this information set should be, by <u>Bayes's formula</u>,

$$\mu_i(x|t_i) = P(x)/\sum_{y\in t_i} P(y),$$

that is, the prior probability of x divided by the sum of prior probabilities of all nodes in the information set t_i, whenever this formula is well-defined (not 0/0).

A <u>belief system</u> μ is a vector that specifies a belief-probability distribution $\mu_i(\bullet|t_i)$ over the nodes of each information set t_i of each player i in the game.

Bayes's formula yields a unique belief system for any full-support behavioral strategy profile.

But for strategy profiles that do not have full support, Bayes's formula may leave some belief probabilities undefined.

A beliefs system μ is <u>consistent</u> with a behavioral strategy profile σ iff there exists a sequence of full-support behavioral strategies $\hat{\sigma}^k$ that converge to σ (all $\hat{\sigma}^k{}_i(d_i|t_i) \to \sigma_i(d_i|t_i)$) and yield Bayesian beliefs $\hat{\mu}^k$ that converge to μ as $k\to\infty$ (all $\hat{\mu}^k{}_i(x|t_i) \to \mu_i(x|t_i)$).

A behavioral-strategy profile σ is <u>sequentially rational</u> given a beliefs system μ iff, at each information set t_i of each player i, $\sigma_i(\bullet|t_i)$ assigns positive move-probabilities only to moves that maximize i's expected payoff at t_i, given i's beliefs $\mu_i(\bullet|t_i)$ about the current node in the information set t_i and given what the behavioral-strategy profile σ specifies about players' behavior after this information set.

A <u>sequential equilibrium</u> is a pair (σ,μ), where σ is a behavioral strategy profile and μ is a belief system, such that σ is sequentially rational given the beliefs system μ, and the beliefs system μ is consistent with the behavioral-strategy profile σ.

The Holdup Problem Player 1 can invest to improve an asset which he may later sell player 2. First player 1 chooses an amount $e \geq 0$ to spend on improving the asset. With this investment, the asset will be worth $v_1(e) = e^{0.5}$ to player 1, but it will be worth $v_2(e) = 2e^{0.5}$ 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 amount $e \geq 0$ to spend on improving the asset.
Player 2 observes this investment e.
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:
$u_1(e, p, \text{accept}) = p - e$, $u_2(e, p, \text{accept}) = v_2(e) - p$, $u_1(e, p, \text{reject}) = v_1(e) - e$, $u_2(e, p, \text{reject}) = 0$.

There is a unique subgame-perfect equilibrium.
At the last stage, player 1 accepts if $p > v_1(e)$ and rejects if $p < v_1(e)$.
So player 2's optimal offer, given e, must be to offer $p = v_1(e)$, which player 1 must accept.
(Note: Player 1 is actually indifferent between accepting and rejecting, but there would be no optimal offer for 2 if player 1 rejected in this case of indifference!)
So player 1 knows that his payoff from e will be $v_1(e) - e = e^{0.5} - e$, which is maximized by $e=0.25$
So the equilibrium outcome is: 1 chooses $e = 0.25$, 2 offers $p = 0.25^{0.5} = 0.5$,
and payoffs are $u_1 = 0.5 - 0.25 = 0.25$, $u_2 = 2 \times 0.25^{0.5} - 0.5 = 1 - 0.5 = 0.5$.

Seller-offer game. First player 1 chooses his investment $e \geq 0$.
Then player 1 chooses the price $p \geq 0$ at which he offers to sell the asset.
Player 2 observes e and p, and then can choose to accept or reject 1's offer. Payoffs are still
$u_1(e, p, \text{accept}) = p - e$, $u_2(e, p, \text{accept}) = v_2(e) - p$, $u_1(e, p, \text{reject}) = v_1(e) - e$, $u_2(e, p, \text{reject}) = 0$.

In the unique subgame-perfect equilibrium, player 2 accepts if $p \leq v_2(e)$ but rejects if $p > v_2(e)$, so given e, player 1 offers $p = v_2(e)$. So player 1 chooses $e = 1$ to maximize $2e^{0.5} - e$.
So the equilibrium outcome is: 1 chooses $e = 1$ and offers $p = 2 \times 1^{0.5} = 2$,
and payoffs are $u_1 = 2 - 1 = 1$, $u_2 = 2 \times 1^{0.5} - 2 = 0$.
This seller-offer game also has many other Nash equilibria that are not subgame perfect.

Notice that the equilibrium sum of payoffs $u_1 + u_2$ is greater in the seller-offer game.
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 $(\hat{e},\hat{p})$ such that $v_2(\hat{e}) \geq \hat{p} \geq \hat{e} + \max_e (v_1(e) - e) = \hat{e} + 0.25$ (such as $\hat{e}=1$, $\hat{p}=1.625$), so that each does better than he could alone. With either player offering the price, there is a Nash equilibrium in which 1 invests this $\hat{e}$, and then this price $\hat{p}$ is offered and accepted, but rejection would follow any other investment $e \neq \hat{e}$ or any other price-offer $p \neq \hat{p}$. These Nash equilibria violate sequential rationality, however, as threats to reject prices between $v_1(e)$ and $v_2(e)$ would not be credible.

Introduction to repeated games Players 1 and 2 will meet on $\tau+1$ days, numbered $0,1,2,...,\tau$.
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
f_1	5, 0	2, 2

(Prisoners' dilemma)

except on the last day τ their payoffs will be:

Player 1: \ Player 2:	g_2	f_2
g_1	5, 5	0, 4
f_1	4, 0	2, 2

(Trust game)

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 $V_i = \sum_{k=0}^{\tau} \delta^k u_{ik}$
for some given discount factor δ between 0 and 1.

If the first payoff matrix (the prisoners' dilemma) were played once, (f_1,f_2) would be the unique equilibrium, yielding the Pareto-dominated payoff allocation (2,2).
But in multi-period games, opportunities to respond later can enlarge the set of equilibria.

For example, consider the strategy "choose g_i until f_1 or f_2 is chosen, and thereafter choose f_i."
We can show that it is an equilibrium here for both players to choose this strategy, if $\delta \geq 2/3$.
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	$3+5\delta$, $3+5\delta$	$0+2\delta$, $5+2\delta$
f_1	$5+2\delta$, $0+2\delta$	$2+2\delta$, $2+2\delta$

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 τ of repetitions of the prisoners' dilemma.
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:

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

(We use $w+w\delta+w\delta^2+...+w\delta^{s-1} = w(1-\delta^s)/(1-\delta)$.)
If $1>\delta\geq 2/3$ then $3(1-\delta^\tau)/(1-\delta)+5\delta^\tau \geq 5+2\delta(1-\delta^\tau)/(1-\delta)$.
As $\tau\rightarrow\infty$, overall payoffs in this equilibrium depend on first-day actions as follows:

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

The War-of-Attrition game. There are two players, numbered 1 and 2, who can meet on day 0, day 1, ... through day τ, to try to get a valuable prize that is worth V.
On each day, if the game has not yet ended, each player can choose to fight or quit.
The game ends as soon as somebody quits, or it ends after day τ if nobody quits.
On each day when both choose to fight, they both lose \$1. On any day when one player fights and the other player quits, the fighter gets the prize worth \$V (and the game ends). If they both quit on the same day, or if they both fight on all days (0 through τ), then nobody gets the prize.

The normal-form strategy for each player i can be described by a number c_i chosen from the set $\{0,1,...,\tau+1\}$, where c_i represents the day when player i would quit, if the other player does not quit first, except that $c_i=\tau+1$ represents the strategy "never quit". So the payoffs functions are
$u_1(c_1,c_2) = V-c_2$ if $c_1 > c_2$, but $u_1(c_1,c_2) = -c_1$ if $c_1 \le c_2$;
$u_2(c_1,c_2) = V-c_1$ if $c_2 > c_1$, but $u_2(c_1,c_2) = -c_2$ if $c_2 \le c_1$.

Suppose player 2 chooses $\tilde{c}_2$ randomly, according to a probability distribution $\sigma_2(t) = P(\tilde{c}_2=t)$.
Player 1's expected payoff from choosing $c_1=d$ is $Eu_1(d,\tilde{c}_2) = \sum_{t<d} (V-t)\sigma_2(t) + \sum_{t\ge d} (-d)\sigma_2(t)$.
So $Eu_1(0,\tilde{c}_2) = 0$, and $Eu_1(d+1,\tilde{c}_2) - Eu_1(d,\tilde{c}_2) = V\sigma_2(d) - \sum_{t>d} \sigma_2(t) = (V+1)\sigma_2(d) - \sum_{t\ge d} \sigma_2(t)$.
(After d days, 1's willingness to fight one more day would earn \$V if $\tilde{c}_2=d$, or lose \$1 if $\tilde{c}_2>d$.)

There is a symmetric full-support randomized equilibrium in which each player i chooses $\tilde{c}_i$ randomly according to a probability distribution $\sigma_1 = \sigma_2$. We can find this σ_2 by solving the equations $0 = Eu_1(d+1,\tilde{c}_2) - Eu_1(d,\tilde{c}_2) = (V+1)\sigma_2(d) - \sum_{t\ge d} \sigma_2(t)$ for all d in $\{0,1,...,\tau\}$.
First, using $\sum_{t\ge 0} \sigma_2(t) = 1$, we get $\sigma_2(0) = 1/(V+1)$.
Then $\sigma_2(1) = [\sum_{t\ge 1} \sigma_2(t)]/(V+1) = [1-\sigma_2(0)]/(V+1) =[1 - 1/(V+1)]/(V+1)$.
Then for all $d=1,...,\tau$, we can recursively compute
$\sigma_2(d) = [\sum_{t\ge d} \sigma_2(t)]/(V+1) = [1-\sum_{t<d} \sigma_2(t)]/(V+1) = [1- 1/(V+1)]^d/(V+1)$.
At the end, we have $\sigma_2(\tau+1) = 1-\sum_{t<\tau+1} \sigma_2(t)$ (which goes to 0 as $\tau\to\infty$).
On each day $d\le\tau$, we have $\sigma_2(d)/\sum_{t\ge d} \sigma_2(t) = 1/(V+1)$. This ratio is the conditional probability of player 2 quitting on day d, given that she has not quit earlier.
So this mixed strategy corresponds to a behavioral strategy in which, on any given day, if nobody has quit earlier, then the probability of player i quitting today is $q = 1/(V+1)$.
This conditional probability q satisfies the equation $qV+(1-q)(-1) = 0$, which makes the other player just indifferent between quitting immediately and fighting one more day.
In this symmetric randomized equilibrium, each player is willing to quit on day 0, and so each player's expected payoff is $0 = Eu_1 = Eu_2$.

There is also a nonsymmetric equilibrium in which player 1 is always expected to fight and player 2 is expected to quit immediately, so that $c_1=\tau+1$, $c_2=0$, $u_1=V$, and $u_2=0$.
There is also a nonsymmetric equilibrium in which player 2 is always expected to fight and player 1 is expected to quit immediately, so that $c_1=0$, $c_2=\tau+1$, $u_1=0$, and $u_2=V$.
These nonsymmetric equilibria can be interpreted as a model of property rights.

Infinitely Repeated games
Infinitely repeated games can be used as simple 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 c_{ik} in some set C_i. In period k, each player i's payoff u_{ik} will depend on both players' actions according to some utility function $u_i:C_1 \times C_2 \to \mathbb{R}$; that is, $u_{ik} = u_i(c_{1k}, c_{2k})$.
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 sum of player i's payoffs is
$V(u_{i1},u_{i2},u_{i3},...) = u_{i1} + \delta u_{i2} + \delta^2 u_{i3} + ... + \delta^{k-1} u_{ik} + ...$
For a constant payoff x each period, the δ-discounted sum would be $x/(1-\delta)$.
The objective of each player i in the repeated game is to maximize the expected discounted sum of his payoffs, with respect to some discount factor δ, where $0 < \delta < 1$.

Fact. (Recursion formula) $V(u_{i1},u_{i2},u_{i3},...) = u_{i1} + \delta V(u_{i2},u_{i3},u_{i4},...)$.

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 a characteristic 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 sum 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: $V_i(\theta) = E[u_i(\tilde{s}_1(\theta),\tilde{s}_2(\theta))) + \delta\, V_i(\Theta(\tilde{s}_1(\theta),\tilde{s}_2(\theta);\theta))]$.

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,
$V_1(\theta) \geq E[u_i(a_1,\tilde{s}_2(\theta))) + \delta\, V_i(\Theta(c_1,\tilde{s}_2(\theta);\theta))]$, for all c_1 in C_1,
$V_2(\theta) \geq E[u_i(\tilde{s}_1(\theta),c_2)) + \delta\, V_i(\Theta(\tilde{s}_1(\theta),c_2;\theta))]$, for all c_2 in C_2.

Example. 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	5, 5	0, 6
d_1	6, 0	1, 1

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 values for the players in the states satisfy the equations:
$V_1(0) = u_1(c_1,c_2) + \delta V_1(0)$, $V_1(1) = u_1(d_1,d_2) + \delta V_1(1)$,
$V_2(0) = u_2(c_1,c_2) + \delta V_2(0)$, $V_2(1) = u_2(d_1,d_2) + \delta V_2(1)$.
So $V_1(0) = 5 + \delta V_1(0)$, $V_1(1) = 1 + \delta V_1(1)$, and so $V_1(0) = 5/(1-\delta)$, $V_1(1) = 1/(1-\delta)$.
Similarly, $V_2(0) = 5/(1-\delta)$, $V_2(1) = 1/(1-\delta)$.
For this scenario to be an equilbrium, we need:
$V_1(0) \geq u_1(d_1,c_2) + \delta V_1(1)$, $V_1(1) \geq u_1(c_1,d_2) + \delta V_1(1)$,
$V_2(0) \geq u_2(c_1,d_2) + \delta V_2(1)$, $V_2(1) \geq u_2(d_1,c_2) + \delta V_2(1)$.
That is, we need: $5/(1-\delta) \geq 6 + \delta 1/(1-\delta)$ and $1/(1-\delta) \geq 0 + \delta 1/(1-\delta)$, which are satisfied when $1 \geq \delta \geq 1/5$.

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 values $V_1(\theta)$ for player 1 in each state θ satisfy the equations:
$V_1(0) = u_1(c_1,c_2) + \delta V_1(0)$, $V_1(1) = u_1(c_1,d_2) + \delta V_1(0)$,
$V_1(2) = u_1(d_1,c_2) + \delta V_1(0)$.
So $V_1(0) = 5 + \delta V_1(0)$, and $V_1(0) = 5/(1-\delta)$.
So $V_1(1) = 0 + \delta 5/(1-\delta)$, and $V_1(1) = 5\delta/(1-\delta)$.
So $V_1(2) = 6 + \delta 5/(1-\delta) = (6-\delta)/(1-\delta)$.
Similarly, $V_2(0) = 5$, $V_2(1) = (6-\delta)/(1-\delta)$, $V_2(2) = 5\delta$.
To have an equilibrium, we need: $V_1(0) \geq u_1(d_1,c_2) + \delta V_1(1)$,
$V_1(1) \geq u_1(d_1,d_2) + \delta V_1(1)$, $V_1(2) \geq u_1(c_1,c_2) + \delta V_1(2)$,
and similar conditions for player 2.
These inequalities become $5/(1-\delta) \geq 6 + \delta 5\delta/(1-\delta)$, $5\delta/(1-\delta) \geq 1 + \delta 5\delta/(1-\delta)$,
$(6-\delta)/(1-\delta) \geq 5+\delta(6-\delta)/(1-\delta)$, which are satisfied when $1 \geq \delta \geq 1/5$.
(Algebraic fact used: $(1-\delta^2) = (1-\delta)(1+\delta)$.)

Duality in linear programming (LP) (MWG appendix M.M, Myerson 1991 pp 125-127.)
Suppose that we are given m×n matrix $A = (a_{ij})_{i\in\{1,...,m\},j\in\{1,...n\}}$,
and vectors $b = (b_i)_{i\in\{1,...,m\}} \in \mathbb{R}^m$ and $c = (c_j)_{i\in\{1,...,n\}} \in \mathbb{R}^n$.
Consider the primal linear-programming problem:

choose x in $\mathbb{R}^n$ to minimize $c'x$ subject to $Ax \geq b$.

This problem is equivalent to: $\text{minimize}_x \max \{c'x + y'(b-Ax) \mid y\in\mathbb{R}^m, y\geq\mathbf{0}\}$,
because if x violated constraint the constraints then b−Ax would have some positive components and so the max here would become +∞ (very bad when we are minimizing).
If we reversed the order of min and max, we would get

choose y in $\mathbb{R}^m$ to maximize $\min\{y'b + (c'-y'A)x \mid x\in\mathbb{R}^n\}$ subject to $y \geq \mathbf{0}$.

This problem is (similarly) equivalent to the dual linear-programming problem:

choose y in $\mathbb{R}^m$ to maximize $y'b$ subject to $y'A = c'$ and $y \geq \mathbf{0}$.

Duality Theorem of Linear Programming. Suppose that the constraints of the primal and dual LP problems both have feasible solutions. Then these problems have optimal solutions x and y such that $c'x = y'b$ (equal values) and $y'(b-Ax) = 0$ (complementary slackness).

Proof If x and y satisfy the primal and dual constraints then we must have

$$c'x \geq c'x + y'(b-Ax) = y'b + (c'-y'A)x = y'b. \qquad (*)$$

So both the primal and dual problems must have bounded optimal values.
Dual boundedness implies that we cannot find any $\hat{y}$ such that $\hat{y}'A=\mathbf{0}'$, $\hat{y}\geq\mathbf{0}$, and $\hat{y}'b>0$,
because otherwise we could infinitely improve any dual solution by adding multiples of $\hat{y}$.
Now for any number θ, linear duality implies that exactly one of the following is true:

(1) $\exists x \in \mathbb{R}^n$ such that $Ax \geq b$ and $-c'x \geq -\theta$.

(2) $\exists(y,\omega) \in \mathbb{R}^m \times \mathbb{R}$ such that $y\geq\mathbf{0}$, $\omega\geq 0$, $y'A - \omega c' = \mathbf{0}$, and $y'b - \omega\theta > 0$.

But when (2) is true, we must have ω>0, or else we would have the vector $\hat{y}$ described above.
So we could divide any solution of (2) through by ω>0 to get a solution of (2) with ω=1.
So whenever (2) holds, we must also have

(2′) $\exists y \in \mathbb{R}^m$ such that $y \geq \mathbf{0}$, $y'A = c'$, and $y'b > \theta$.

So the dual maximization can do better than any value θ that is below the minimal value of the primal. Thus, the optimal values of the primal and dual problems must be equal.
It can be shown that the set of feasible values for each problem is a closed set, and so optimal solutions actually exist. At optimal solutions x and y, we have $y'b=c'x$,
which (by (*) above) implies the complementary slackness equation $y'(b-Ax) = 0$. QED.

[Our basic linear duality thm can also be seen as a special case of duality in linear programming.
Given the matrix A and the vector b from basic linear duality, suppose we let $c=\mathbf{0}$ in $\mathbb{R}^n$.
Then the constraints of the dual LP problem always have a feasible solution $y=\mathbf{0}$.
So LP duality implies that this primal is feasible ($\exists x$ s.t. $Ax\geq b$) if and only if
its and its dual share the same optimal value, which must be $c'x = \mathbf{0}'x = 0$,
and so any $y\geq\mathbf{0}$ with $y'A=\mathbf{0}$ must have $y'b\leq 0$.]