Exercise 8.9.1

(a) The strategic form can be represented by two tri-matrices. In every field, the lower left hand payoff represents the payoff to player I and the upper right hand entry is the payoff to II. The middle entry is the payoff of player III. Player I is the row player, II is the column player, and player III chooses between the left and the right tri-matrix.

b_{III}

	b	c
b	$\frac{1}{4}$ (upper right), $\frac{1}{4}$ (middle), $\frac{1}{4}$ (lower left)	1 (upper right), $\frac{1}{3}$ (middle), $\frac{1}{3}$ (lower left)
c	$\frac{1}{3}$ (upper right), $\frac{1}{3}$ (middle), 1 (lower left)	

c_{III}

	b	c
b	$\frac{1}{3}$ (upper right), 1 (middle), $\frac{1}{3}$ (lower left)	$\frac{1}{2}$ (upper right), $\frac{1}{2}$ (middle), $\frac{1}{2}$ (lower left)
c	$\frac{1}{2}$ (upper right), $\frac{1}{2}$ (middle), $\frac{1}{2}$ (lower left)	$\frac{1}{3}$ (upper right), $\frac{1}{3}$ (middle), $\frac{1}{3}$ (lower left)

(b) The game is symmetric. Take any permutation π of $\{1, 2, 3\}$, and it is immediate to check that $u_i(a_1, a_2, a_3) = u_{\pi(i)}\left(a_{\pi(1)}, a_{\pi(2)}, a_{\pi(3)}\right)$ for any player i and any strategy profile (a_1, a_2, a_3).

(c) No.

Exercise 8.9.2

The reaction curves cross nowhere else. See figure 8.9.2.

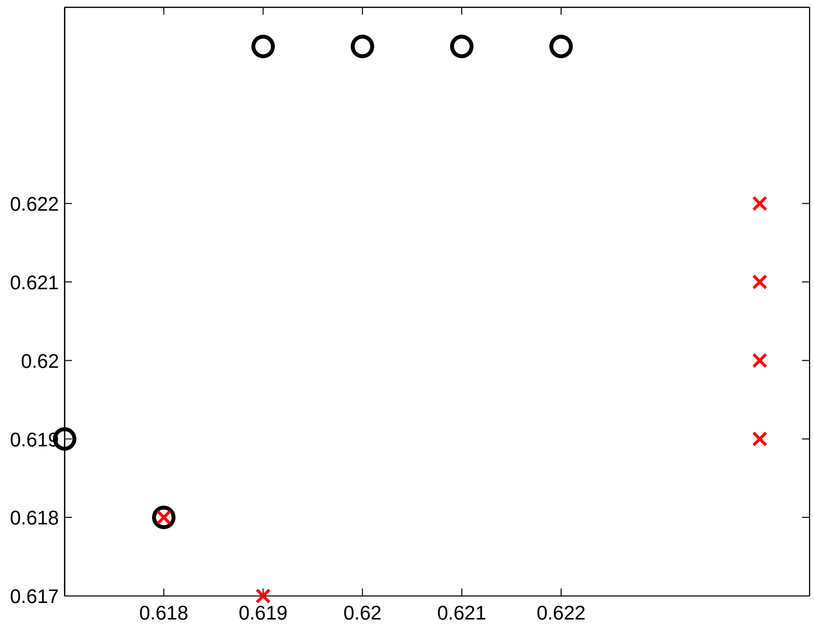

Figure 8.9.2

Exercise 8.9.3

See figure 8.9.3.

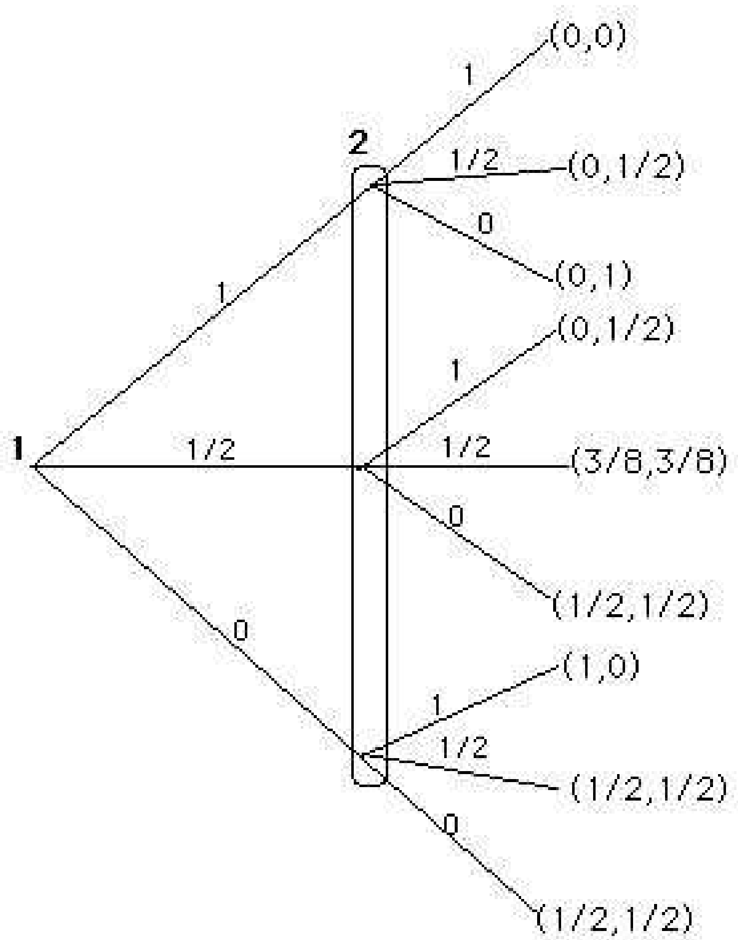

Figure 8.9.3

Exercise 8.9.4

Now $\pi(d, e) = (1 - d)(1 - e)$. Clearly, the unique Nash equilibrium (in fact in dominant strategies) is $d = e = 0$.

Exercise 8.9.5

If a strongly dominated strategy s were played with a positive probability by player I, then he could profitably deviate to playing the strategy that dominates s, contradicting the fact that s was a Nash equilibrium strategy. The following game has a weakly dominated Nash equilibrium (D, R).

	L	R
U	1, 1	0, 0
D	0, 0	0, 0

There always exists a Nash equilibrium where no weakly dominated strategy is used with a positive probability. First, eliminate all weakly dominated strategies of the game G to obtain the game G'. In G' there exists a Nash equilibrium by the existence theorem. So denote by (p', q') the equilibrium strategies of the players in G'. They are not weakly dominated in G' by construction. We need to show that they constitute an equilibrium in G and that they are not weakly dominated in G. The second one is obvious since all the strategies that are in G and not in G' are those that are weakly dominated by some strategies in G'. Hence, no strategy in G can weakly dominate anything in G'. Now assume that (p', q') weren't an equilibrium of G. Then there would exist a strategy p', say for I, such that $(p', q) \succ (p, q)$ but this means that p' isn't weakly dominated (since it dominates something in G'), hence, p' is a strategy in G', which is a contradiction.

Exercise 8.9.6

Take player I. A completely mixed strategy for him is one where $Aq = \alpha e$. Because the matrix A is non-singular, this equation has a unique solution for each α. Since we also require q to be a probability measure, we get precisely one α^* which satisfies all these conditions. Similarly for player II.

Exercise 8.9.7

To show the first inequality observe that $\forall q, p\ \Pi_1(p,q) \leq \max_{p\in P} \Pi_1(p,q)$ thus for every p we have $\min_{q\in Q} \Pi_1(p,q) \leq \min_{q\in Q} \max_{p\in P} \Pi_1 p, q$. Since this inequality holds for evey $p \in P$, it also has to hold for the maximum, which is what we needed to show. To see that the second inequality holds, notice that at the Nash equilibrium strategies $(\tilde{p}, \tilde{q})$ both players are indifferent between any of the strategies that they choose with positive probability, and those strategies yield at least as much as the strategies that are not played at all. In particular, when $\tilde{q}$ is played by player II, it holds for player I's payoff that $\Pi_1(\tilde{p}, \tilde{q}) \geq \max_{p\in P} \Pi_1(p, \tilde{q}) \geq \min_{q\in Q} \max_{p\in P} \Pi_1(p,q)$. The last inequality is clear: the minimum will always yield at most as much as if we picked a particular point.

Exercise 8.9.8

Let p be the probability of that I plays *up* and q the probability that II plays left. To compute the mixed Nash equilibrium we have that II has to be indifferent between his two strategies, given I's randomization, which yields $2p + 4 - 4p = 5p + 3 - 3p$ or $4 - 2p = 3 + 2p$, so that $p = 1/4$. Similarly, q has to be such that I is indifferent between his two strategies, which yields $2q+2 = 4q+1$, or $q = 1/2$. The payoffs to the players are $(3, \frac{7}{2})$. To obtain II's security strategy we compute his payoffs under each of I's strategies, if II randomizes with q, that is, $E_{up}(q) = 5 - 3q$, $E_{down}(q) = 3 + q$. II's security strategy is thus $q = \frac{1}{2}$, with the same payoff as in Nash equlibrium.

Exercise 8.9.9

Without any preplay communication or a device to break symmetry there is no way for the players to coordinate on either of the pure-strategy equilibria. The mixed Nash equilibrium seems the only possible candidate because it is symmetric in the way it treats players, and each player is then indifferent between both of his strategies, so that mixing is a best reply. To compute the mixed Nash equilibrium, let p be probability player I assigns to playing *box* and let q be the probability II assigns to *box*. By the indifference condition we get $2 - 2p = p$ and $1 - q = 2q$, so that $p = \frac{2}{3}, q = \frac{1}{3}$. To compute the security strategies we compute II's payoffs, given q, for each of I's actions, to get $E_{box}(q) = q, E_{ball}(q) = 2 - 2q$, so that $q = \frac{2}{3}$. Similarly we obtain $p = \frac{1}{3}$. If II plays his security strategy and I plays mixed Nash then I obtains a payoff 1, rather than $\frac{2}{3}$ which is what he obtains in the mixed Nash or as his security payoff. Hence, the reason players do not switch to their security strategies is that it would not be a best reply for at least one of them to do so if the other one did.

8.9.10

As in exercise 8.9.9 we obtain that in a mixed Nash equilibrium player I plays *box* with probability $\frac{2}{3}$ and player II plays *box* with probability $\frac{1}{3}$. Similarly, we get that their security strategies require them to reverse these probabilities. However, if II now plays his security strategy while I plays his mixed Nash strategy, then I gets -1 instead of $-\frac{2}{3}$ which is his security payoff.

Exercise 8.9.11

The risk dominant equilibrium is (s_1, t_1) and the Pareto-dominant equilibrium is (s_2, t_2).

Exercise 8.9.12

The following game satisfies the desired conditions:

	L	R
U	1, 1	1, −1
D	−1, 1	2, 2

In general, we could denote the payoffs to player I by $\begin{matrix} a & b \\ c & d \end{matrix}$, symmetrically for player II, and derive that in order for equilibrium (U,L) to be risk-dominant, we need the following inequality to hold: $\frac{1}{2} \geq \frac{b-d}{(a-c+b-d)}$.

Exercise 8.9.13

Another possible equilibrium is that money is completely worthless, and since then no one can possibly buy anything with the money, it is completely worthless. So money being worth something is just one of the possible conventions.

Exercise 8.9.14

Clearly, in Boston, the equilibrium where the gentleman goes first, should be the focal one. In some other, more pleasant city, perhaps the other equilibrium may be the focal one.

Exercise 8.9.15

It seems that the red sector between the two green ones should be the focal one, because it is the only one that doesn't have a symmetric counterpart. Again, ths is a matter of taste.

Exercise 8.9.16

The first Why follows from the fact that y_k is the profit maximizing output for p_k. The second Why is problematic, because it assumes that we can find such a convergent sequence y_k. Notice that $y_k = s(p_k)$ doesn't tell us anything about either boundedness or convergence of the sequence y_k. If we knew that $s(.)$ was continuous then indeed y_k would have to be bounded, since p_k is such, and a convergent subsequence of y_k would exist, so we could apply the arguments in the exercise to that subsequence. But this line of proof doesn't work, since we are proving that $s(.)$ is continuous! We need to be more careful. Here is a very careful proof. Assume that all components of p are non-zero. Take the point $y = s(p)$ and first show that y_k is a bounded sequence. Suppose this isn't the case, so there exists a subsequence y_{k_l} that goes to infinity (more precisely, the $i-$th component $y^i_{k_l}$ of y_{k_l} goes to infinity, as $l \to \infty$, for some i). But then also $p^T_{k_l} y_{k_l} \to \infty$, and since $p_{k_l} \to p$, this would contradict the fact that y was optimal production at price p. Thus, the sequence y_k has to be bounded, so it must have at least one convergent subsequence. Now we will show that every convergent subsequence of y_k has to converge to y. Take a convergent subsequence y_{k_l} of y_k and assume it converges to $\bar{y} \neq y$. Then, $p^T y > p^T \bar{y}$ - because of the strict convexity, the profit-maximizing y is unique. Thus, there is an $\varepsilon > 0$ so that $p^T y > p^T \bar{y} + \varepsilon$. Also, since the scalar product of two vectors is a continuous function, there exist a $\delta > 0$ such that $\left|\tilde{p}^T \tilde{y} - p^T \bar{y}\right| < \frac{\varepsilon}{2}$ and $\left|\tilde{p}^T y - p^T y\right| < \frac{\varepsilon}{2}$ whenever $||\tilde{y} - \bar{y}|| < \delta$ and $||\tilde{p} - p|| < \delta$. Since p is the limit of p_{k_l} and $\bar{y}$ is the limit of y_{k_l} there also exist l_1 and l_2 so that $||y_{k_l} - \bar{y}|| < \delta$ for every $l \geq l_1$ and $||p_l - p|| < \delta$ for every $l > l_2$. Define $l^* = \max\{l_1, l_2\}$. Then, for every $l > l^*$ $\left|p^T_{k_l} y_{k_l} - p^T \bar{y}\right| < \frac{\varepsilon}{2}$. On the other hand $\left|p^T_{k_l} y - p^T y\right| < \frac{\varepsilon}{2}$, so that $p^T_{k_l} y \geq p^T y - \frac{\varepsilon}{2} > p^T \bar{y} + \frac{\varepsilon}{2} > p^T_{k_l} y_{k_l}$, which contradicts the fact that y_{k_l} was the optimal production at price p_{k_l}. Hence every convergent subsequence of y_k converges to y, so the sequence y_k is in fact convergent, and y is its limit, implying that $s(.)$ is continuous. Notice, that we need the fact that no component of p is equal to 0 - otherwise y_k could be an unbounded sequence.

Exercise 8.9.17

(a) Let $X = S_1 \times S_2 \times ... \times S_n$ and let $F = (G_1, G_2, ..., G_n)$. For each i, assume that S_i is convex and compact, and that G_i is non-empty, convex valued, and the graph of G_i is closed in $X \times S_i$. Then, X is convex and compact, and the graph of F is convex valued, non-empty and has a closed graph in $X \times X$. Thus, by Kakutani's theorem, there exists an $\tilde{s} \in X$ such that $\tilde{s} \in G(\tilde{s})$, or equivalently, $\tilde{s}_i \in G_i(\tilde{s}_i)$.

(b) Now denote $M_i(s) = \arg\max_{s_i \in G_i(s)} u_i(s)$. If $u_i(s_{-i}, s_i)$ is a quasi-concave function of s_i for every s_{-i}, then M_i is non-empty valued and convex correspondence with a closed graph, $M_i : X \to S_i$. Now let $H = (M_1, M_2, ..., M_n)$. Again, by Kakutani's theorem there exists an $s^* \in H(s^*)$, which is the optimal choice for each player i.

Exercise 8.9.18

For instance: Act only on the maxim on which everyone will be willing to act.

Exercise 8.9.19

(a) In a two-party context, if the median voter votes for a party, then either all the voters to his left or all the voters to his right vote for that same party. Thus, whichever of the two parties is able to capture the vote of the median voter, will get the majority of votes.

(b) The expected position of the median voter is $\frac{1}{2}$, so that if a party positions itself at say $c > \frac{1}{2}$, and if the other party locates at $\frac{1}{2}$, it will be able to win with a very high probability (if we assume that the number of voters is very large, then this probability is close to 1).

(c) Suppose that the Idealists and the Formalists locate at $\frac{1}{2}$. Then, if the Intuitionists locate at $\frac{1}{2} - \varepsilon$, then they will capture $\frac{1-\varepsilon}{2}$ votes, while the other two parties only get a half of the remaining vote each, that is $\frac{1+\varepsilon}{4}$ each, so the Intuitionists would win. So, by a similar argument, the other two parties must position themselves at symmetric points around $\frac{1}{2}$. Suppose they place their platforms at c and $1-c$, where $c > \frac{1}{4}$. Then, the Intuitionists should place their platform at $c - \varepsilon$ and get $c - \frac{\varepsilon}{2} > \frac{1}{4}$ votes as opposed to getting $\frac{1-2c}{2} < \frac{1}{4}$ votes if they locate at $\frac{1}{2}$. But then the other two parties could have done better if they had placed at $c - \frac{\varepsilon}{2}$ so that this could not be an equilibrium. When $c = \frac{1}{4}$, it becomes optimal for the Intuitionists to place at $\frac{1}{2}$, and the two incumbents have no incentives to place their platforms at a lower c.