SETTLED EQUILIBRIA

by Roger Myerson and Jörgen Weibull

very preliminary notes, Nov 16, 2011

Concepts and terminology here are tentative and may change in later versions!

http://home.uchicago.edu/~rmyerson/research/settled.pdf

Example 1. The "battle of sexes" game

	a_2	b_2
a_1	3, 2	0, 0
b_1	0, 0	2, 3

Three equilibria: (a_1, a_2), (b_1, b_2), $(0.6[a_1]+0.4[b_1], 0.4[a_2]+0.6[b_2])$.
If this game is played only once by players who have no cultural or historical context, the symmetric mixed equilibrium might be an appropriate prediction. But in culturally familiar settings, we may expect players to develop understandings that coordinate their expectations at (a_1,a_2) or (b_1,b_2).

How to formalize this intuition?
The totally mixed equilibrium is perfect and proper and K-M stable...

Familiar game are played in a historical context where players know how similar players have played similar games.
In such histories, norms regularly develop that specify which alternatives the players should (and should not) consider.

Institutions develop when people stop considering alternatives that are physically available.
Such an institution is self-enforcing if, when people play rationally within institutional norms, no player could do better by violating these norms.

Given a finite strategic-form game $G = (N, S, (u_i)_{i \in N})$ with players $N = \{1,2,...n\}$, strategies $S = \times_{i \in N} S_i$, each S_i a nonempty finite set, utility functions $u_i:S \to \mathbb{R}$.

$\Delta(S_i)$ is the set of probability distributions over S_i.
A profile of mixed strategies is any $\sigma \in \times_{i \in N} \Delta(S_i)$.
We extend $u_i:\times_{i \in N} \Delta(S_i) \to \mathbb{R}$ by $u_i(\sigma) = \sum_{s \in S} (\prod_{i \in N} \sigma_i(s_i))\, u_i(s)$.
Let $u_i(\sigma_{-i},[s_i])$ denote the expected payoff that player i would get from choosing pure strategy s_i when everyone else randomizes according to σ.

A Nash equilibrium is any σ in $\times_{i \in N} \Delta(S_i)$ such that
$u_i(\sigma) \geq u_i(\sigma_{-i},[s_i])\ \forall s_i \in S_i,\ \forall i \in N$.

Def. A block is any $T = \times_{i \in N} T_i$ where $\emptyset \neq T_i \subseteq S_i$ $\forall i \in N$.
(Represents pure strategies that are "expected" under some social norm.)

A <u>curb set</u> (Basu-Weibull, 1991) is any block T such that, for every mixed strategy profile σ with support in T ($\times_{i \in N} \{s_i | \sigma_i(s_i)>0\} \subseteq T_i$ $\forall i \in N$), all best replies are in T ($\text{argmax}_{s_i \in S_i} u_i(\sigma_{-i},[s_i]) \subseteq T_i$).
Curb sets may be viewed as basic models of self-enforcing institutions.

The good equilibria of Example 1 are in minimal curb sets of the game.

Example 2:

	a_2	b_2
a_1	1, 1	1, 1
b_1	1, 1	0, 0

Minimal curb sets do not exclude the dominated equilibrium here.

A related concept yields minimal blocks with no weakly dominated strategies:
A block is <u>absorbing</u> (Kalai-Samet 1984) iff there is some ε>0 such that, for every mixed strategy profile that puts probability at least 1−ε on the block, every player has at least one best reply in the block.

Example 3.

	a_2	b_2	c_2	d_2
a_1	0, 5	1, 4	0, 3	1, 0
b_1	1, 0	0, 3	1, 4	0, 5

Equilibria: $(0.75[a_1]+0.25[b_1], 0.5[a_2]+0.5[b_2])$,
$(0.5[a_1]+0.5[b_1], 0.5[b_2]+0.5[c_2])$,
$(0.25[a_1]+0.75[b_1], 0.5[c_2]+0.5[d_2])$.

The only curb set or absorbing block includes the whole game.
What strategic refinement can exclude the middle equilibrium?

The support of the middle equilibrium is a subgame that has two other pure equilibria, which fail to be equilibria of the given game.
The supports of the left and right equilibria, as subgames, contain no other subgame equilibria.

If rational responses within an accepted block quickly lead towards its equilibria, then, to be self-enforcing, the block only needs to be absorbing near its equilibria.

Def. Given block T, player i, and mixed strategy profile σ, let $\beta_i^T(\sigma)$ be the set of ρ_i in $\Delta(S_i)$ such that $\{r_i \in S_i | \rho_i(r_i)>0\} \subseteq \text{argmax}_{t_i \in T_i} u_i(\sigma_{-i},[t_i])\}$ (best replies in T_i).
Def. σ is an equilibrium of the restricted game on T iff $\sigma_i \in \beta_i^T(\sigma)$ $\forall i \in N$.
Def. A block T is quasi-absorbing iff every equilibrium of the restricted game on T has a neighborhood in $\times_{i \in N} \Delta(S_i)$ where every i's best reply in T_i is optimal in S_i.

Def. For any $\varepsilon>0$ and any block T, an ε-subequilibrium on T is any σ in $\times_{i \in N} \Delta(S_i)$ such that $\sum_{t_i \in T_i} \sigma_i(t_i) > 1-\varepsilon$ and $u_i(\sigma) + \varepsilon > \max_{t_i \in T_i} u_i(\sigma_{-i},[t_i])$.
Fact. T is quasi-absorbing iff there exists some $\varepsilon>0$ such that, for every σ that is an ε-subequilililibrium on T, $\max_{t_i \in T_i} u_i(\sigma_{-i},[t_i]) = \max_{s_i \in S_i} u_i(\sigma_{-i},[s_i])$, $\forall i \in N$.
Proof: If such an $\varepsilon>0$ exists, its ε-subequilibria include neighborhoods around all equilibria in the T-game.
If not, then we have a sequence of ε-subequilibria for which some player i's best reply in T_i is strictly worse than his best reply in S_i, and some convergent subsequence of these must converge to an equilibrium of the T-game as $\varepsilon \to 0$, and thus T is not quasi-absorbing.
Fact. If T is quasi-absorbing then all equilibria of the restricted game on T are Nash equilibria of the original game G.
Fact. The weakly undominated pure strategies of G form a quasi-absorbing set.
Fact. Any quasi-absorbing block T contains the support $\times_{i \in N} \{s_i | \sigma_i(s_i)>0\}$ of at least one proper equilibrium σ of the game G.

Def. An equilibrium σ is settled iff its support is contained in a minimal quasi-absorbing set.
Fact. A settled equilibrium exists.

Example 4. An elaborated version of the battle of sexes game:

	ax_2	ay_2	bx_2	by_2
ax_1	3, 2	3, 2	4,-4	-4, 4
ay_1	3, 2	3, 2	-4, 4	4,-4
bx_1	4,-4	-4, 4	2, 3	2, 3
by_1	-4, 4	4,-4	2, 3	2, 3

Equilibria, for $1/4 \le p \le 3/4$, $1/8 \le q \le 7/8$:
$(p[ax_1]+(1-p)[ay_1],\ q[ax_2]+(1-q)[ay_2])$,
$(q[bx_1]+(1-q)[by_1],\ p[bx_2]+(1-p)[by_2])$,
$(0.3[ax_1]+0.3[ay_1]+0.2[bx_1]+0.2[by_1],\ 0.2[ax_2]+0.2[ay_2]+0.3[bx_2]+0.3[by_2])$.

Any p and q in [0,1] make $(p[ax_1]+(1-p)[ay_1],\ q[ax_2]+(1-q)[ay_2])$ a subgame equilibrium on the block $\{ax_1,ay_1\}\times\{ax_2,ay_2\}$, and best replies from the outer parts of this set lead out of the block.
The only quasi-absorbing set includes all strategies in the game.
To exclude the totally mixed equilibrium, some rational discipline is needed on trembles outside the block, even as norms keep players usually in the block.
Properness...

Def. Player i's <u>ε-proper trembling replies</u> to strategy profile σ is $\beta_i^{\varepsilon}(\sigma)$, the set of mixed strategies ρ_i in $\Delta(S_i)$ such that, $\forall r_i \in S_i$: $\rho_i(r_i)>0$ (totally mixed) and $\rho_i(r_i) \leq \varepsilon\rho_i(s_i)$ for every s_i in S_i such that $u_i(\sigma_{-i},[r_i]) < u_i(\sigma_{-i},[s_i])$.
Def. For any ε>0, a mixed strategy profile σ is <u>ε-proper</u> iff $\sigma_i \in \beta_i^{\varepsilon}(\sigma)$, $\forall i \in N$.
Def. A <u>proper equilibrium</u> is any limit of ε-proper mixed strategy profiles as $\varepsilon \to 0$.

Def. Given ε>0 and block T, a strategy profile σ in $\times_{i \in N} \Delta(S_i)$ is <u>ε-subproper on T</u> iff, for every player i, $\sigma_i \in (1-\varepsilon)\beta_i^{T}(\sigma) + \varepsilon\beta_i^{\varepsilon}(\sigma)$ (probability mixtures of $1-\varepsilon$ times a best reply on T, plus ε times a proper trembling reply).
Fact. $\forall \varepsilon>0$, for any block T, there exist strategy profiles that are ε-subproper in T.
Proof: Consider a modified game in which each player i chooses a strict ordering of his pure strategies in S_i, with the one restriction that the first in the ordering must be a strategy in T_i. Each ordering is interpreted as a mixed strategy of the original game where, $\forall k \in \{1,\ldots,\#S_i\}$, the k'th in the ordering gets probability $\varepsilon^{k-1} / \sum_{h \in \{1,\ldots \#S_i\}} \varepsilon^{h}$. A Nash equilibrium of this modified game is ε-subproper in T.
Fact. Suppose σ is ε-subproper on T, with $0 < \varepsilon < 1/(\max_{i \in N} \#T_i)$.
Then σ is also ε-proper iff $\max_{t_i \in T_i} u_i(\sigma_{-i},[t_i]) = \max_{s_i \in S_i} u_i(\sigma_{-i},[s_i])$, $\forall i \in N$.
Fact. Any limit, as $\varepsilon \to 0$, of strategy profiles that are ε-subproper in T is an equilibrium of the restricted game on T.

Def. A <u>subproper equilibrium on T</u> is any limit as $\varepsilon \to 0$ of strategy profiles that are ε-subproper on T.
Fact. On any block T, there must exist at least one subproper equilibrium.

Def. A block T is <u>subproper</u> iff there exists some $\tilde{\varepsilon}>0$ such that, for every $\varepsilon \in [0,\tilde{\varepsilon}]$, every strategy profile that is ε-subproper on T is also ε-proper for the game G.

Fact. Any quasi-absorbing block is subproper.

Fact. If the block T is subproper then any subproper equilibrium on T is a proper equilibrium of the original game G.

Def. An equilibrium is <u>properly settled</u> iff it is a subproper equilibrium on a minimal subproper block that is a subset of a minimal quasi-absorbing block.

Fact. Properly settled equilibria exist. Any minimal quasi-absorbing block contains at least one properly settled equilibrium.

Example 4. An elaborated version of the battle of sexes game:

	ax_2	ay_2	bx_2	by_2
ax_1	3, 2	3, 2	4,-4	-4, 4
ay_1	3, 2	3, 2	-4, 4	4,-4
bx_1	4,-4	-4, 4	2, 3	2, 3
by_1	-4, 4	4,-4	2, 3	2, 3

For this Example 4, the minimal quasi-absorbing blocks are $\{ax_1,ay_1\}\times\{ax_2,ay_2\}$ and $\{bx_1,by_1\}\times\{bx_2,by_2\}$. There are no smaller subproper blocks, and game has two properly settled equilibria: $(0.5[ax_1]+0.5[ay_1], 0.5[ax_2]+0.5[ay_2])$, and $(0.5[bx_1]+0.5[by_1], 0.5[bx_2]+0.5[by_2])$.

In a totally mixed strategy profile that is ε-subproper on $\{ax_1,ay_1\}\times\{ax_2,ay_2\}$,
if the probability of ax_1 were more than ay_1,
then 2's probability of trembling to bx_2 would be much less than to by_2,
but then player 1 would strictly prefer ay_1 over ax_1.

Example 5. A version of the battle of sexes game where 1 can enter or stay out:

	a_2	b_2
o_1	2,-1	2,-1
ea_1	3, 1	0, 0
eb_1	0, 0	1, 3

Equilibria: $([o_1], p[a_2]+(1-p)[b_2])$ for $0 \le p \le 2/3$, and (eb_1,b_2).
The unique minimal quasi-absorbing set is $\{ea_1\}\times\{a_2\}$.
So the unique settled equilibrium is (ea_1,a_2).

But $\{o_1\}\times\{b_2\}$ is also a subproper block, containing the proper equilibrium (o_1,b_2).

Example 6. Beer-quiche (normal form of game in Kohlberg-Mertens Figure 14).
For 1: bb beer always, bq beer if strong, qb beer if weak, qq quiche always.
For 2: ff fight always, fn fight if beer, nf fight if quiche, nn never fight.

	ff	fn	nf	nn
bb	-21,-8	-21,-8	-1, 0	-1, 0
bq	-2,-8	-18,-9	-2, 1	0, 0
qb	-30,-8	-12, 1	-28,-9	-10, 0
qq	-29,-8	-9, 0	-29,-8	-9, 0

Equilibria: ([bb], (1−p)[nf]+p[nn]) for $0 \le p \le 1/2$,
([qq], (1−p)[fn]+p[nn]) for $0 \le p \le 1/2$, for $0 \le p \le 1/2$.

{bb,bq}×{nf,nn} is quasi-absorbing with the first component of the equilibrium set being its subgame equilibria.
Within this absorbing block, the subblock {bb}×{nf} is subproper, with subproper equilibrium the limit of ε-subproper $((1[bb]+\varepsilon[bq]+\varepsilon^2[qb]+\varepsilon^3[qq])/(1+\varepsilon+\varepsilon^2+\varepsilon^3)$, $(1[nf]+\varepsilon[nn]+\varepsilon^2[ff]+\varepsilon^3[fn])/(1+\varepsilon+\varepsilon^2+\varepsilon^3))$.

But a quasi-absorbing set that includes {qq}×{fn,nn} is must also include bq to exclude (qq,nn) from being a subgame equilibrium, but this adds the subgame equilibrium (bq,nn), which requires nf to be added, and then bb must be added.

Example 7. Another elaborated version of the battle of sexes game:

	ax_2	ay_2	bx_2	by_2
ax_1	3, 2	3, 2	1,-4	-1, 4
ay_1	3, 2	3, 2	-1, 4	1,-4
bx_1	4,-1	-4, 1	2, 3	2, 3
by_1	-4, 1	4,-1	2, 3	2, 3

This game has one minimal quasi-absorbing set $\{bx_1,by_1\}\times\{bx_2,by_2\}$. Its settled equilibria are $(p[bx_1]+(1-p)[by_1], q[bx_2]+(1-q)[by_2])$ for any p and q in [0,1]. Within this minimal quasi-absorbing set, the unique properly settled equilibrium is $(0.5[bx_1]+0.5[by_1], 0.5[bx_2]+0.5[by_2])$.
(But the a-block is also subproper, and on it we find one subproper equilibrium $(0.5[ax_1]+0.5[ay_1], 0.5[ax_2]+0.5[ay_2])$.)