On the Tradeoff Between Efficiency and Strategyproofness

Felix Brandt
(joint work with Haris Aziz, Florian Brandl, and Markus Brill)

GAMES 2016, Maastricht

DSS
Decision Sciences & Systems

Probabilistic Social Choice

- Agents have complete and transitive preference relations $\succcurlyeq_i$ over a finite set of alternatives A.
- A *social decision scheme* f maps a preference profile $(\succcurlyeq_1, \ldots, \succcurlyeq_n)$ to a lottery $\Delta(A)$.

1	**1**	**1**
a	b	a
b	a	c
c	c	b

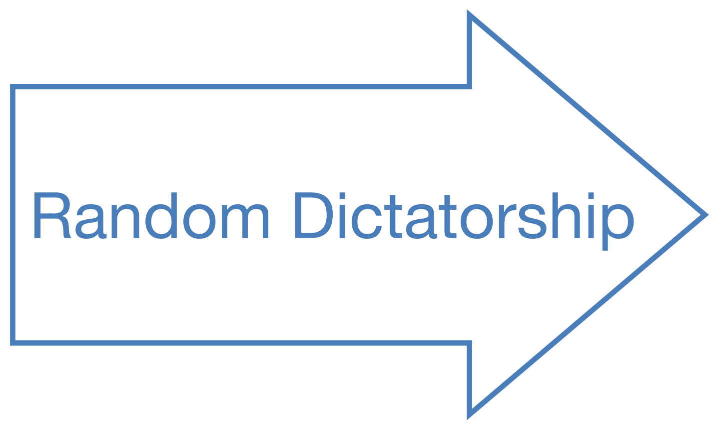

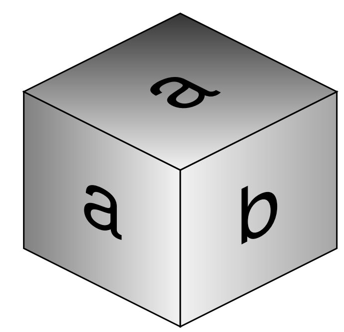

$\frac{2}{3}\, a + \frac{1}{3}\, b$

- Special case: Random assignment (aka house allocation). A is the set of deterministic assignments.
  - Agents are indifferent between all assignments in which they are assigned the same object.

efficiency

No agent can be made better off without making another one worse off

strategyproofness

No agent can obtain a more preferred outcome by misreporting his preferences efficiency

strategyproofness

Only Dictatorship

strict preferences; Gibbard (1973), Satterthwaite (1975)

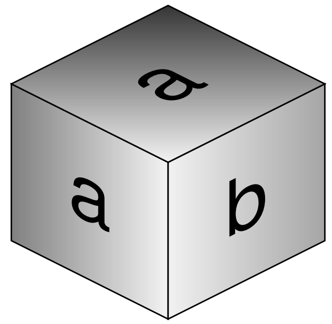

$\frac{2}{3}\, a + \frac{1}{3}\, b$

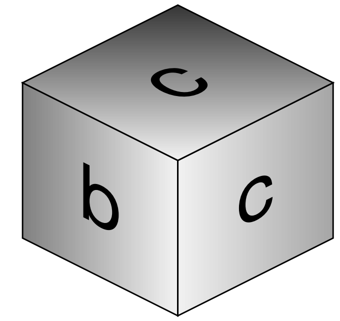

$\frac{1}{3}\, b + \frac{2}{3}\, c$

$a \succ b \succ c$

efficiency

there is no $p \in \Delta(A)$ such that
$p \succcurlyeq_i f(\cdot)$ for all $i \in N$ and
$p \succ_i f(\cdot)$ for some $i \in N$

strategyproofness

there is no $\succcurlyeq_i'$ such that
$f(\succcurlyeq_i', \cdot) \succ_i f(\succcurlyeq_i, \cdot)$

Extend preferences over alternatives to (incomplete) preferences over lotteries!

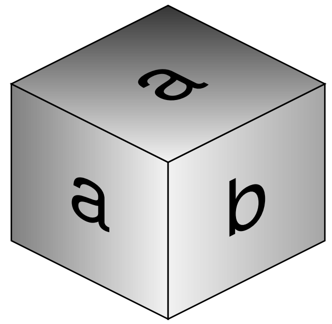

$\frac{2}{3}\, a + \frac{1}{3}\, b$

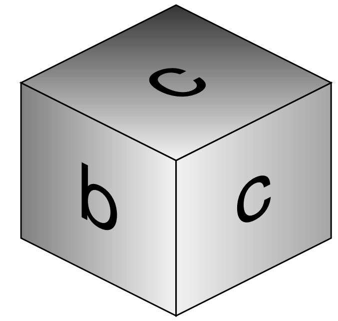

$\frac{1}{3}\, b + \frac{2}{3}\, c$

$a \succ b \succ c$

efficiency

there is no $p \in \Delta(A)$ such that
$p \succcurlyeq_i f(\cdot)$ for all $i \in N$ and
$p \succ_i f(\cdot)$ for some $i \in N$

strategyproofness

weak: there is no $\succcurlyeq_i'$ such that $f(\succcurlyeq_i', \cdot) \succ_i f(\succcurlyeq_i, \cdot)$

strong: for all $\succcurlyeq_i'$ it holds that $f(\succcurlyeq_i, \cdot) \succcurlyeq_i f(\succcurlyeq_i', \cdot)$

Extend preferences over alternatives to (incomplete) preferences over lotteries!

Sure Thing (ST)

	a	$\succ$	b	$\succ$	c
$p =$	(2⁄3		1⁄3		0)
$q =$	(0		1⁄3		2⁄3)

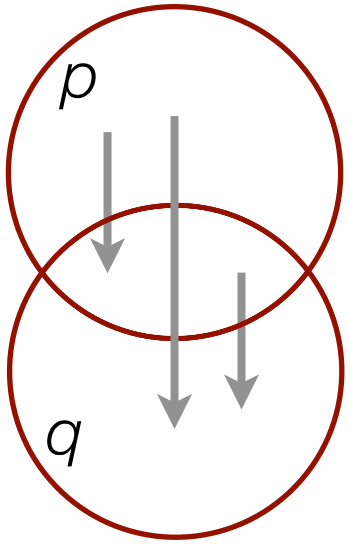

- $p \succcurlyeq^{ST} q \quad \Leftrightarrow \quad \forall x \in supp(p) \setminus supp(q), y \in supp(q): x \succ y$

  $\land \quad \forall x \in supp(p), y \in supp(q) \setminus supp(p): x \succ y$

  $\land \quad \forall x \in supp(p) \cap supp(q): p(x) = q(x)$

  - loosely based on Savage's sure-thing principle
  - inspired by non-probabilistic preference extensions due to Fishburn (1972) and Gärdenfors (1979)

Bilinear Dominance (BD)

	a	$\succ b$	$\succ c$
$p=$	(½	½	0)
$q=$	(⅓	⅓	⅓)

- $p \succcurlyeq^{BD} q \quad \Leftrightarrow \quad [\forall x,y \in A\colon x \succ y \Rightarrow p(x)\, q(y) \geq p(y)\, q(x)]$
  - for every pair of alternatives, it's more likely that *p* yields the better alternative and *q* the worse alternative
  - *p* is preferred to *q* for every consistent SSB utility function
  - Fishburn (1984), Aziz et al. (2015)

$$\forall \succcurlyeq\colon \succcurlyeq^{ST} \subseteq \succcurlyeq^{BD}$$

Stochastic Dominance (SD)

	a	$\succ b$	$\succ c$
$p=$	(½	0	½)
$q=$	(0	½	½)

- $p \succcurlyeq^{SD} q \quad \Leftrightarrow \quad \forall x \in A: \sum_{y \succcurlyeq x} p(y) \geq \sum_{y \succcurlyeq x} q(y)$
  - for every alternative, it's more likely that *p* yields something better
  - *p* yields more expected utility for every consistent vNM function
  - Bogomolnaia & Moulin (2001) and many others

$$\forall \succcurlyeq: \succcurlyeq^{ST} \subseteq \succcurlyeq^{BD} \subseteq \succcurlyeq^{SD}$$

Pairwise Comparison (PC)

	a	$\succ b$	$\succ c$
$p=$	(⅔	0	⅓)
$q=$	(0	1	0)

- $p \succcurlyeq^{PC} q \quad \Leftrightarrow \quad \forall x \in A: \sum_{x \succcurlyeq y} p(x)\, q(y) \geq \sum_{x \succcurlyeq y} q(x)\, p(y)$
  - it's more likely that *p* yields a better alternative
  - minimizes ex ante regret
  - $\succcurlyeq^{PC}$ is a complete relation for all $\succcurlyeq$
  - Blavatskyy (2006), Aziz et al. (2015)

$$\forall \succcurlyeq: \ \succcurlyeq^{ST} \subseteq \succcurlyeq^{BD} \subseteq \succcurlyeq^{SD} \subseteq \succcurlyeq^{PC}$$

weaker

strong
SD-strategyproofness

PC-efficiency

PC-strategyproofness

SD-efficiency

SD-strategyproofness

ex post efficiency

BD-strategyproofness

ST-strategyproofness

Only Random Dictatorship
strict preferences; Gibbard (1977)

weaker

strong *SD*-strategyproofness

PC-efficiency | *PC*-strategyproofness

SD-efficiency | *SD*-strategyproofness

ex post efficiency | *BD*-strategyproofness

ST-strategyproofness

No assignment rule

strict preferences; Bogomolnaia & Moulin (2001)

weaker

	strong *SD*-strategyproofness
PC-efficiency	*PC*-strategyproofness
SD-efficiency	*SD*-strategyproofness
ex post efficiency	*BD*-strategyproofness
	ST-strategyproofness

Probabilistic Serial (PS) assignment rule

strict preferences, Bogomolnaia & Moulin (2001)

weaker

strong
SD-strategyproofness

PC-efficiency — *PC*-strategyproofness

SD-efficiency | *SD*-strategyproofness

ex post efficiency | *BD*-strategyproofness

ST-strategyproofness

No anonymous and neutral social decision scheme
Aziz, Brandl, & B. (2014)

weaker

strong *SD*-strategyproofness

PC-efficiency — *PC*-strategyproofness

SD-efficiency — *SD*-strategyproofness

ex post efficiency — *BD*-strategyproofness

ST-strategyproofness

No anonymous and neutral social decision scheme

Brandl, B., & Geist (2016)

SD Impossibility

- requires at least 4 agents and at least 4 alternatives
  - more than 31 million possible preferences profiles
- was shown with the help of a computer (SMT solver)
- proof has been extracted from the solver's output and brought into human-readable form
- operates on 47 canonical preference profiles and is very tedious to check
- has been verified by a computer (Isabelle/HOL)

SD Impossibility

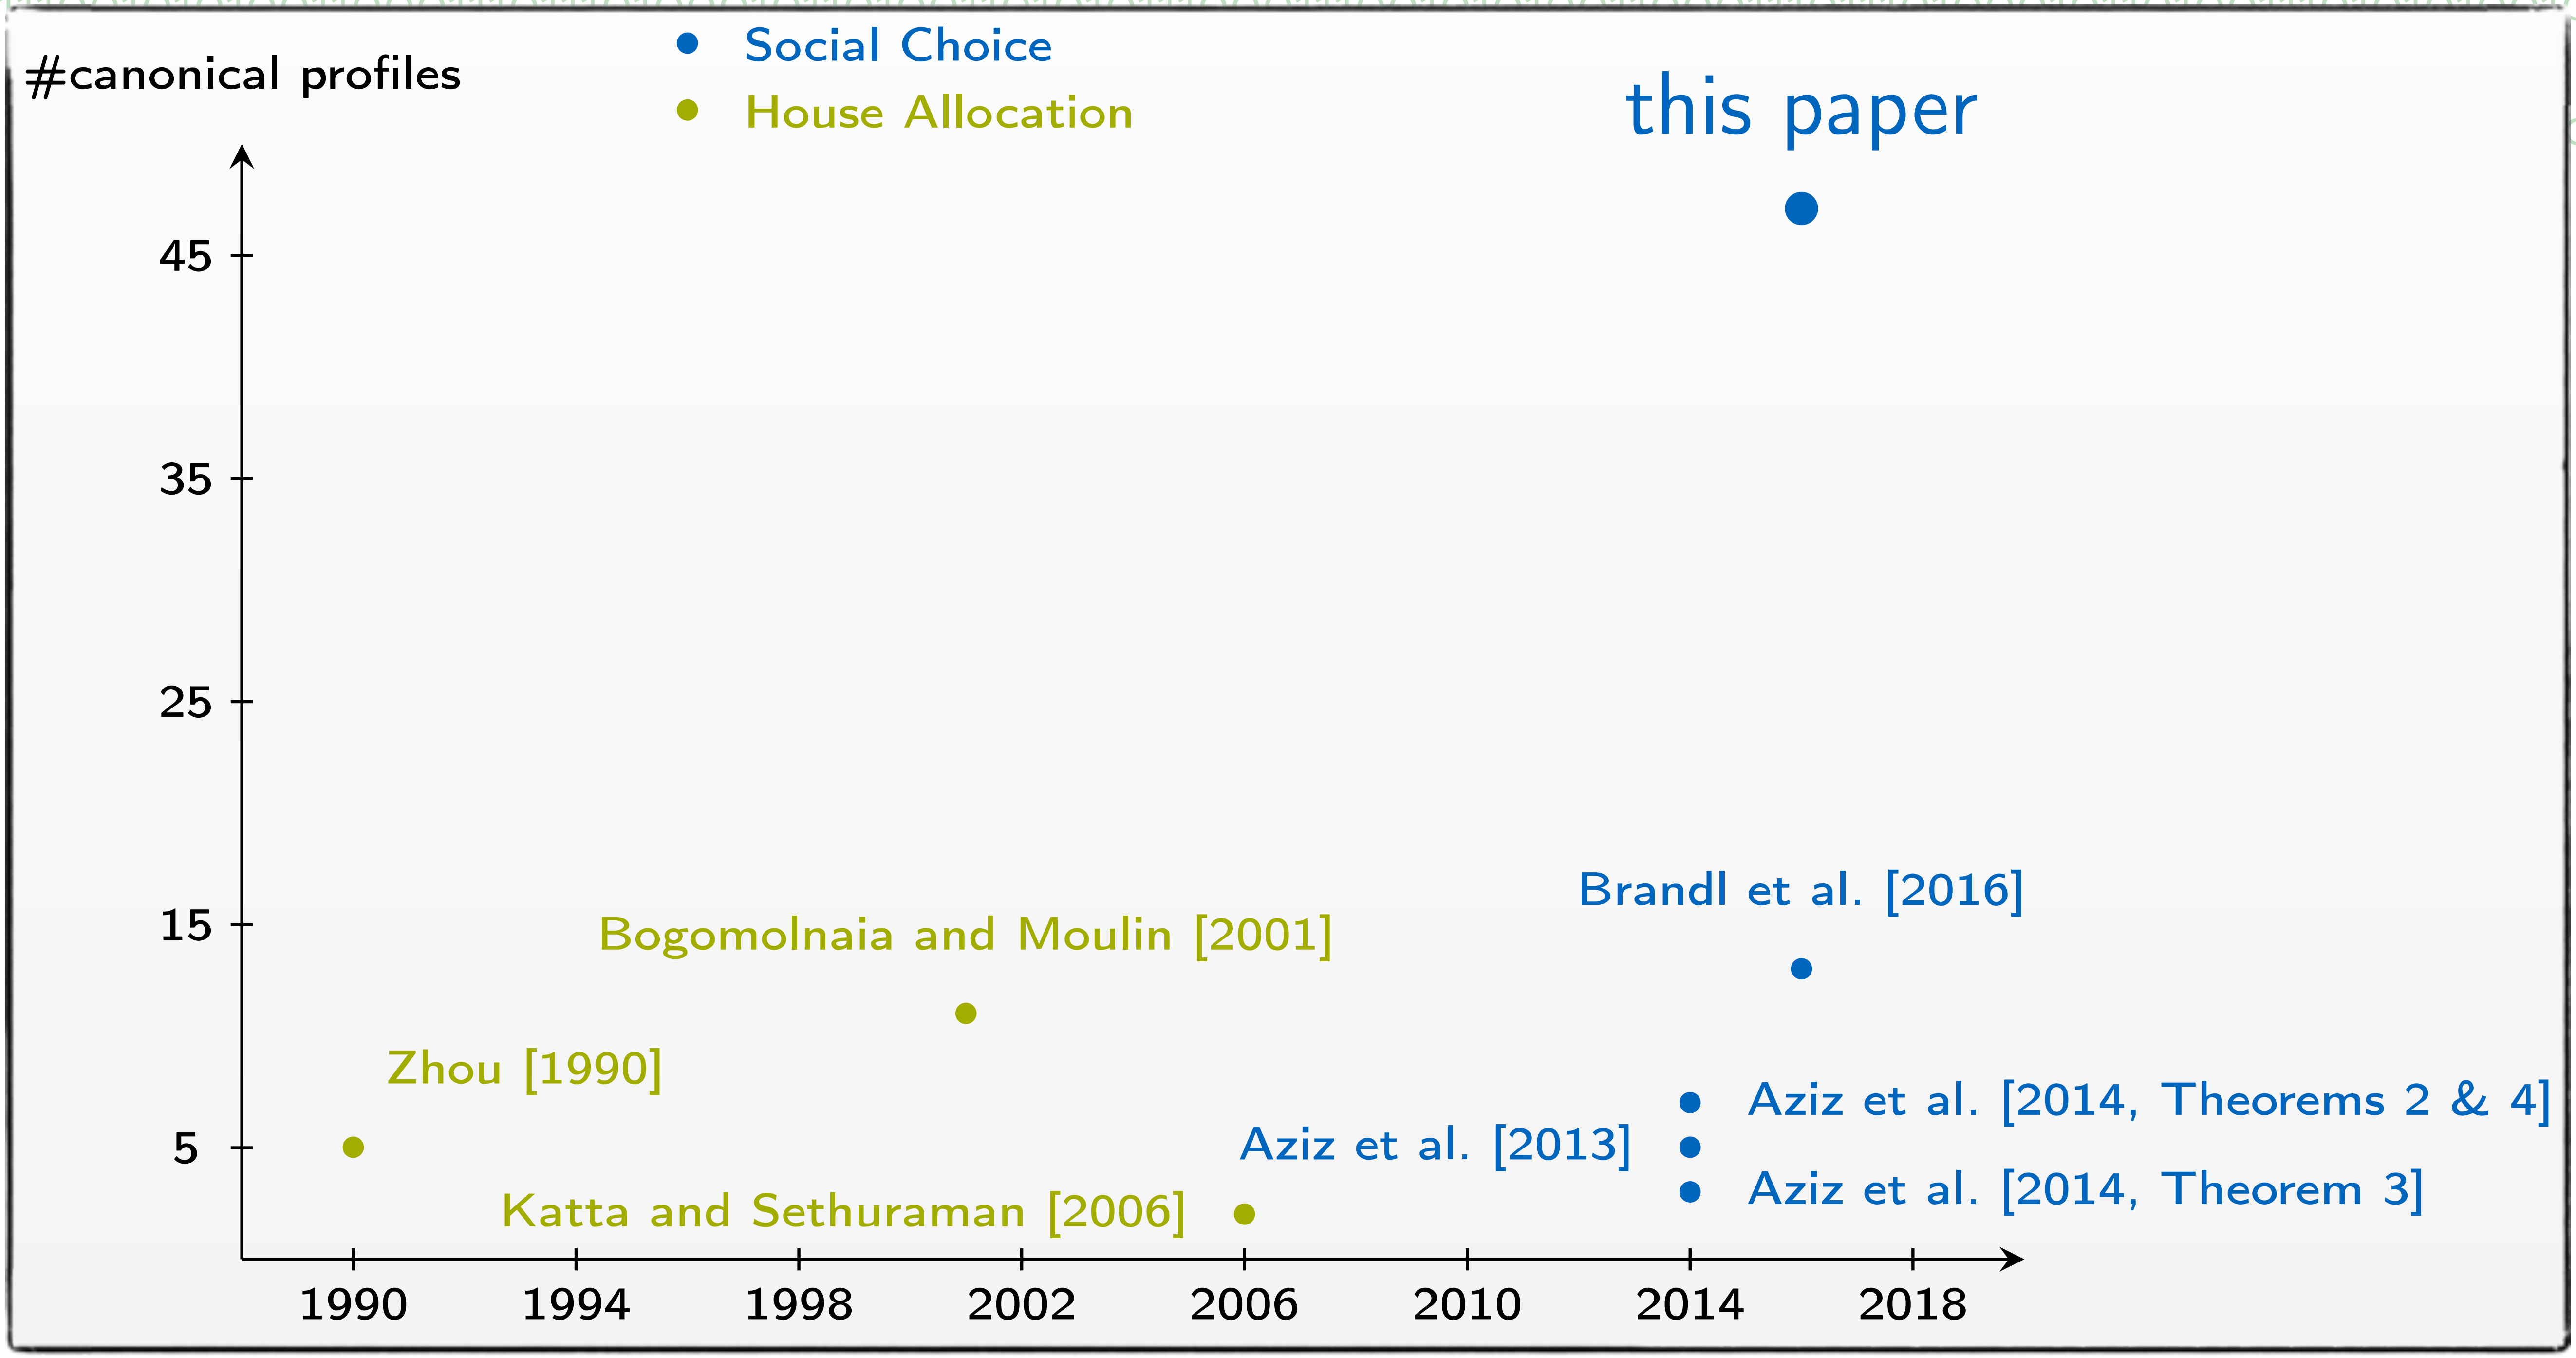

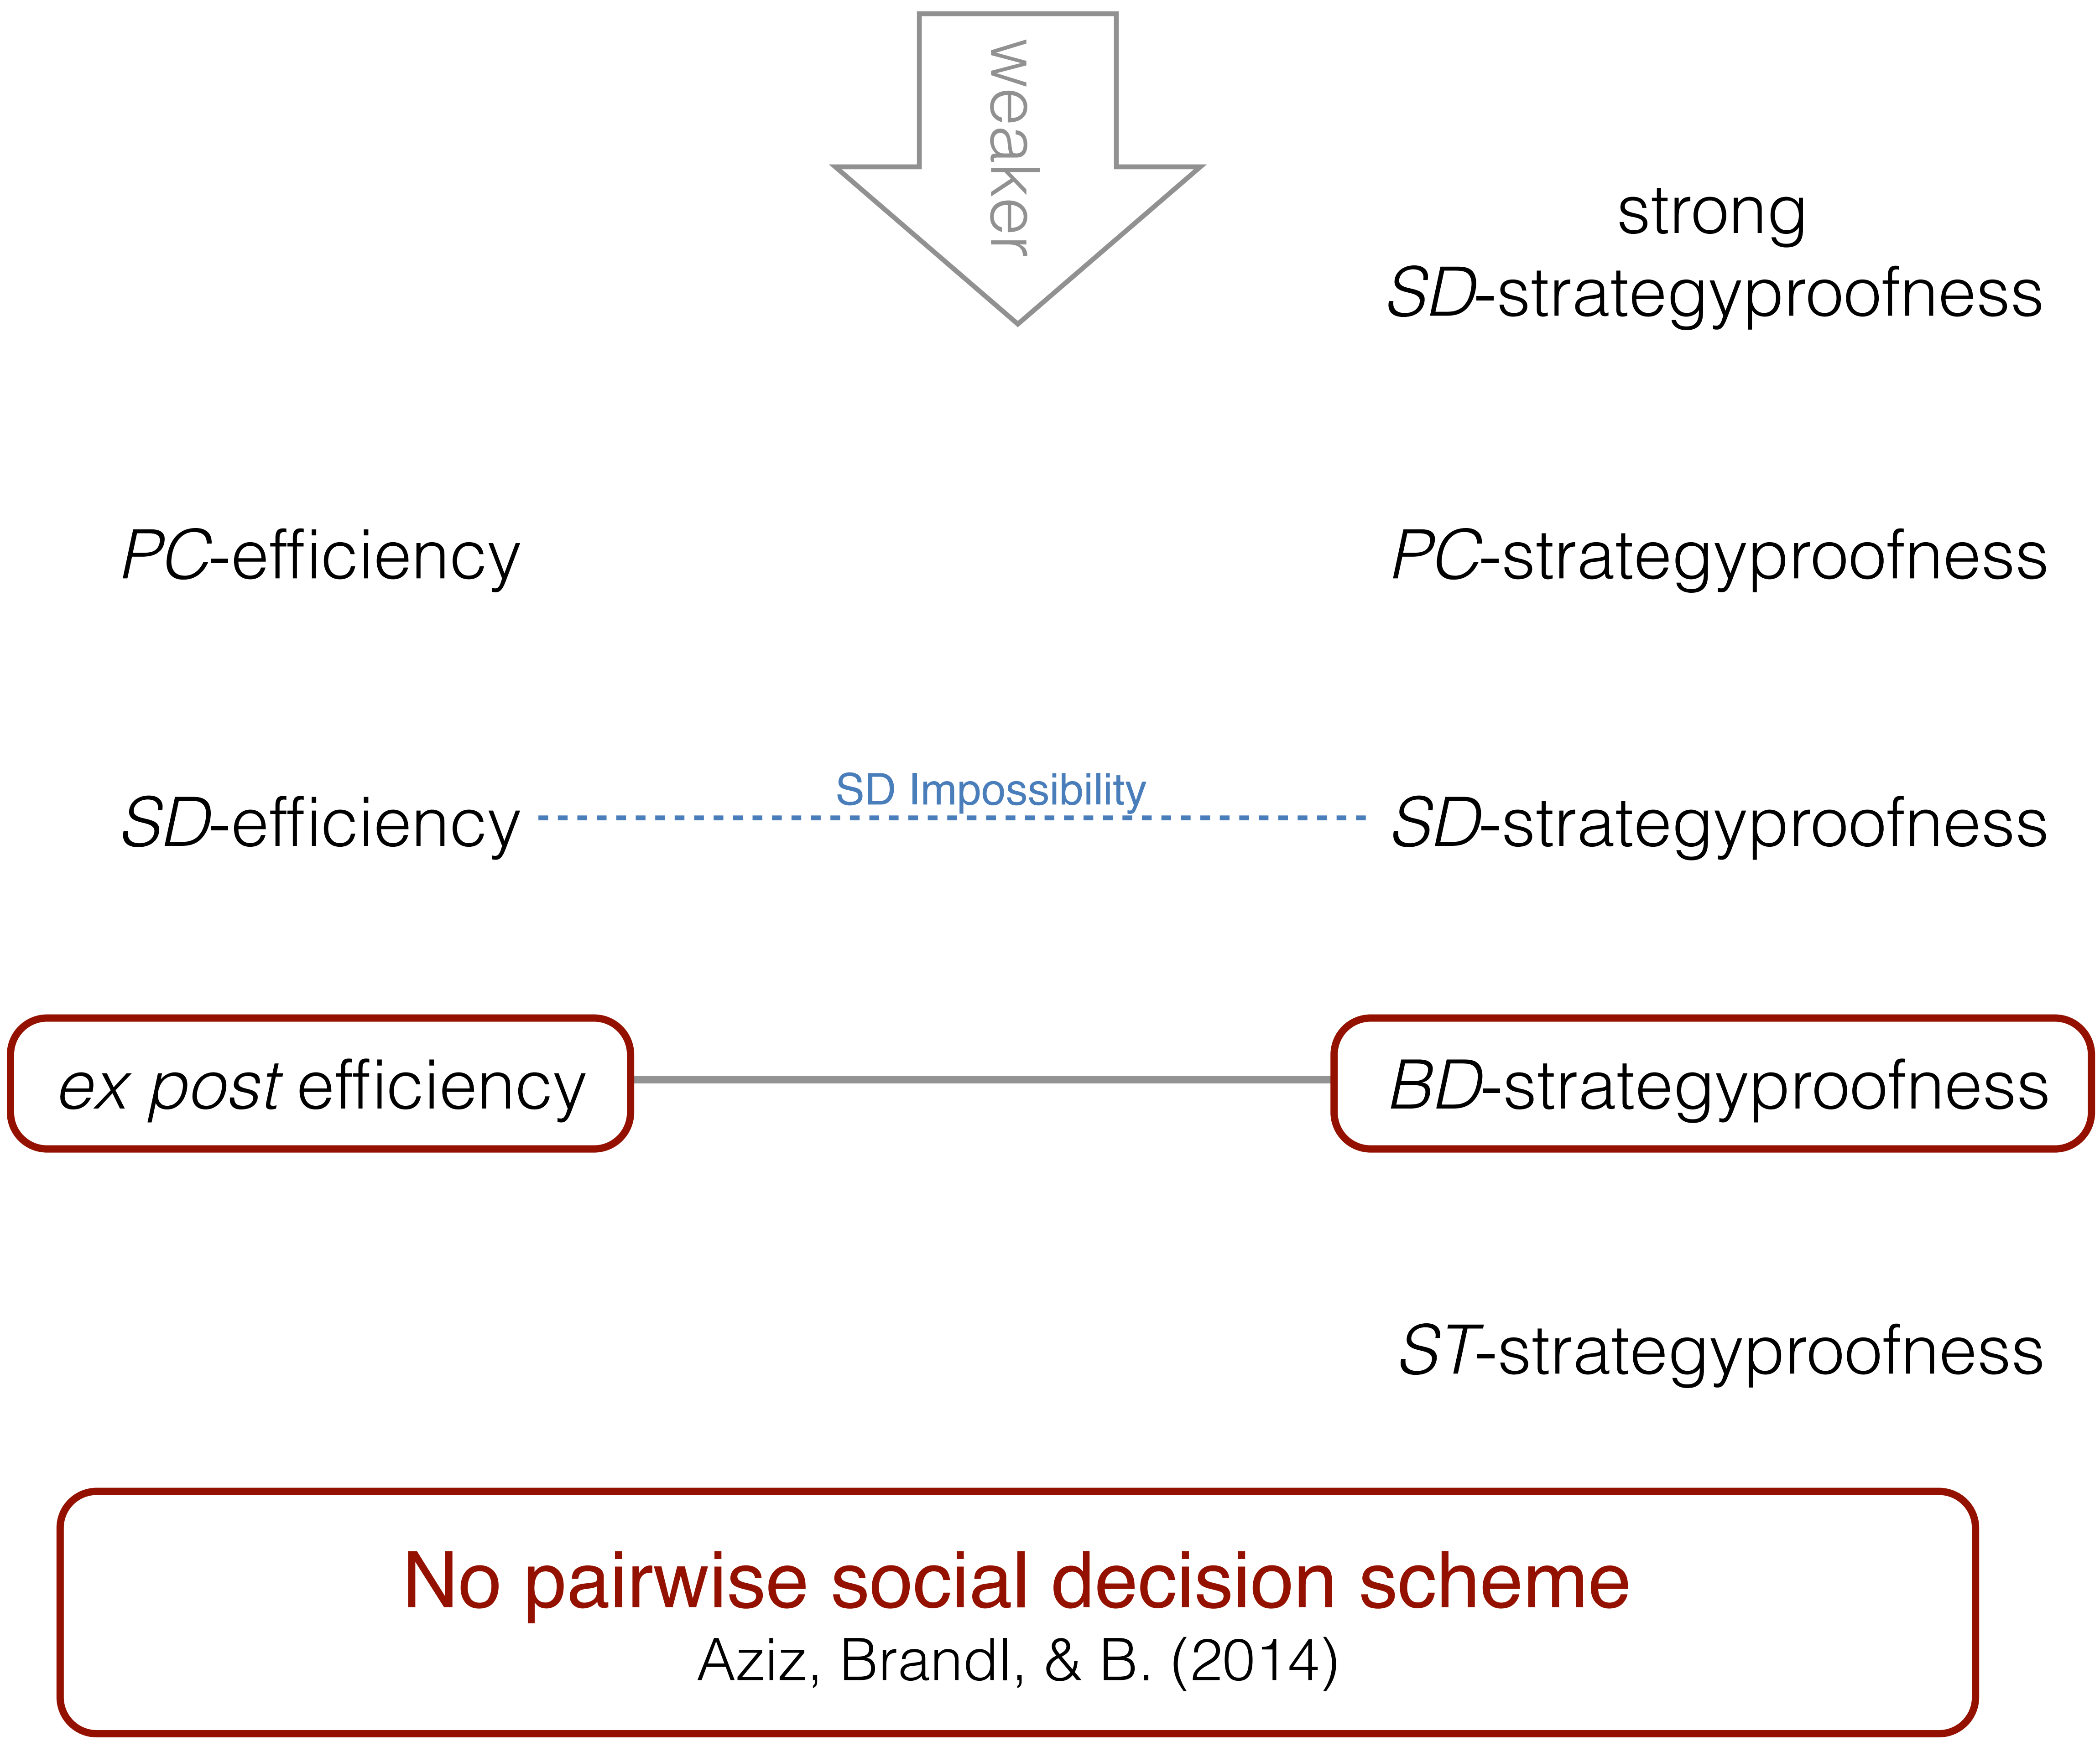
weaker
strong
SD-strategyproofness
PC-efficiency
PC-strategyproofness
SD Impossibility
SD-efficiency
SD-strategyproofness
ex post efficiency
BD-strategyproofness
ST-strategyproofness
No pairwise social decision scheme
Aziz, Brandl, & B. (2014)

Random Serial Dictatorship

- Extension of random dictatorship to weak preferences
  - pick an ordering of agents uniformly at random
  - sequentially narrow down the set of alternatives by letting each agent restrict it to his most preferred ones.
- Widespread assignment rule (aka random priority)

1	1	1
a,c	b,c	a
b	a	b
		c

$\frac{1}{2}\,a + \frac{1}{6}\,b + \frac{1}{3}\,c$

1,2,3: c
1,3,2: a
2,1,3: c
2,3,1: b
3,1,2: a
3,2,1: a

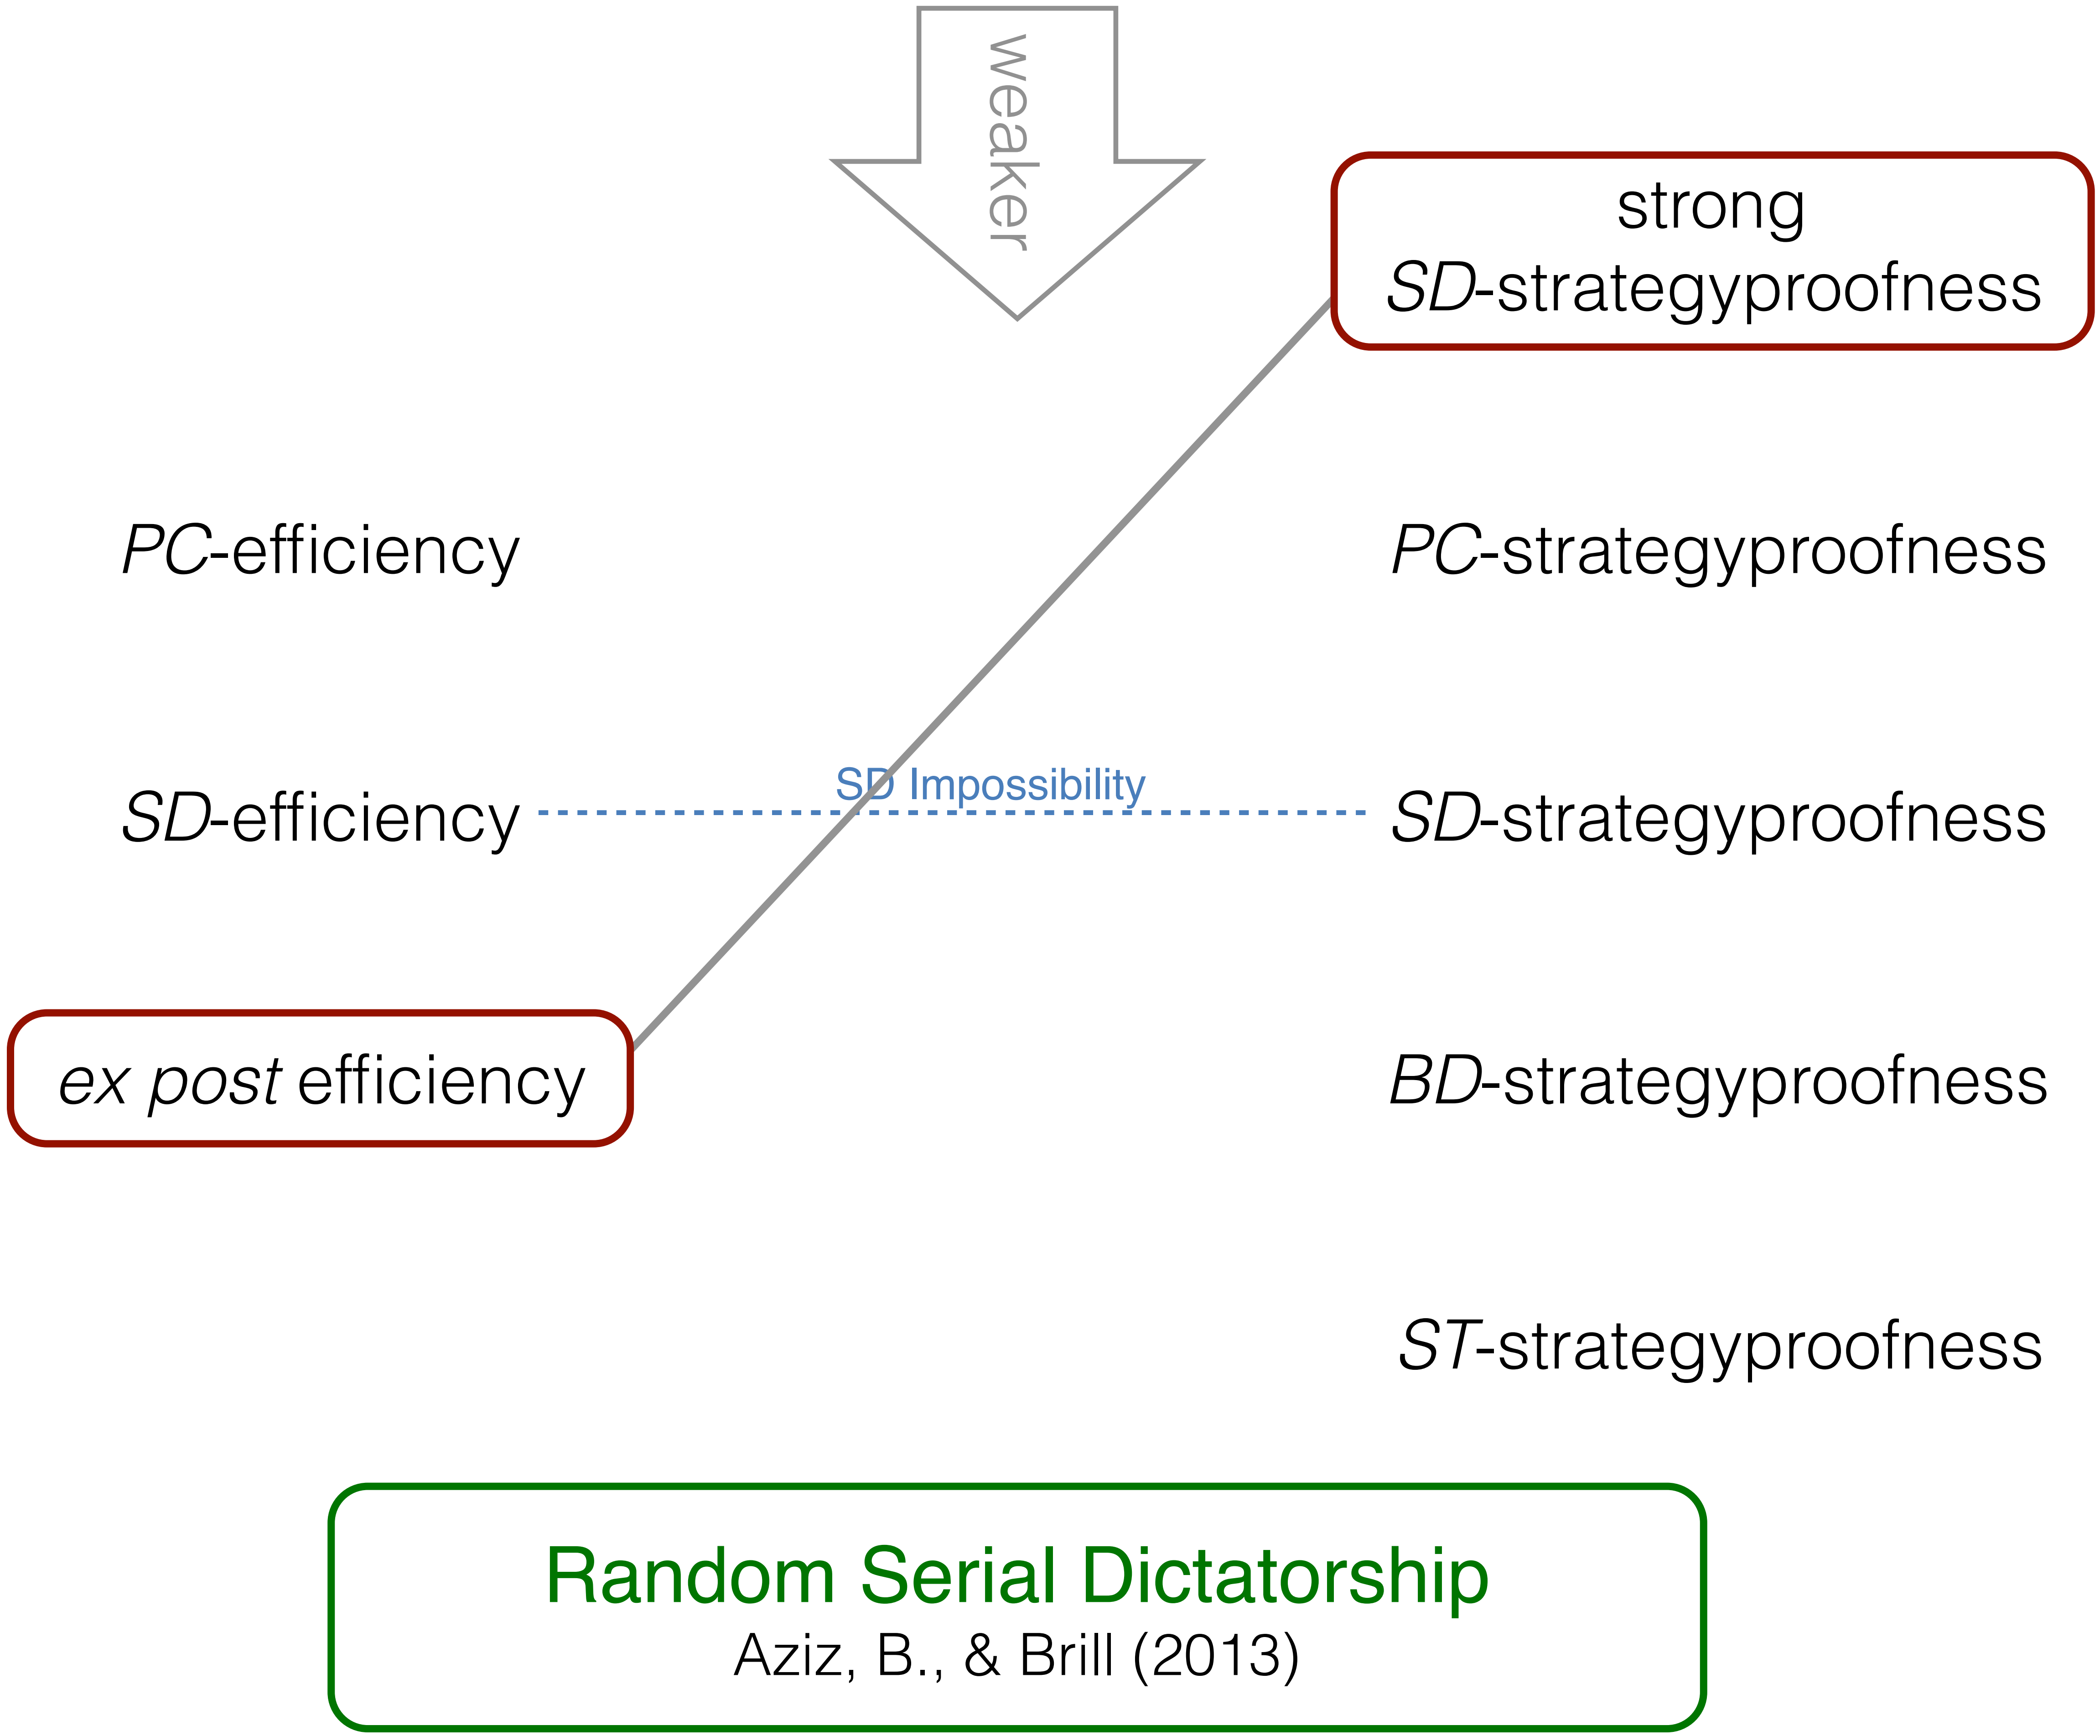
weaker
strong
SD-strategyproofness
PC-efficiency
PC-strategyproofness
SD Impossibility
SD-efficiency
SD-strategyproofness
ex post efficiency
BD-strategyproofness
ST-strategyproofness
Random Serial Dictatorship
Aziz, B., & Brill (2013)

Germain Kreweras

Maximal Lotteries

Peter C. Fishburn

- First studied by Kreweras (1965) and Fishburn (1984)
  - rediscovered by Laffond et al. (1993), Felsenthal and Machover (1992), Fisher and Ryan (1995), Rivest and Shen (2010)
  - preference profiles induce symmetric zero-sum games
  - maximal lotteries correspond to mixed maximin strategies in these games

2	**2**	**1**
a	*b*	*c*
b	*c*	*a*
c	*a*	*b*

	a	*b*	*c*
a	0	1	-1
b	-1	0	3
c	1	-3	0

$\frac{3}{5}\, a + \frac{1}{5}\, b + \frac{1}{5}\, c$

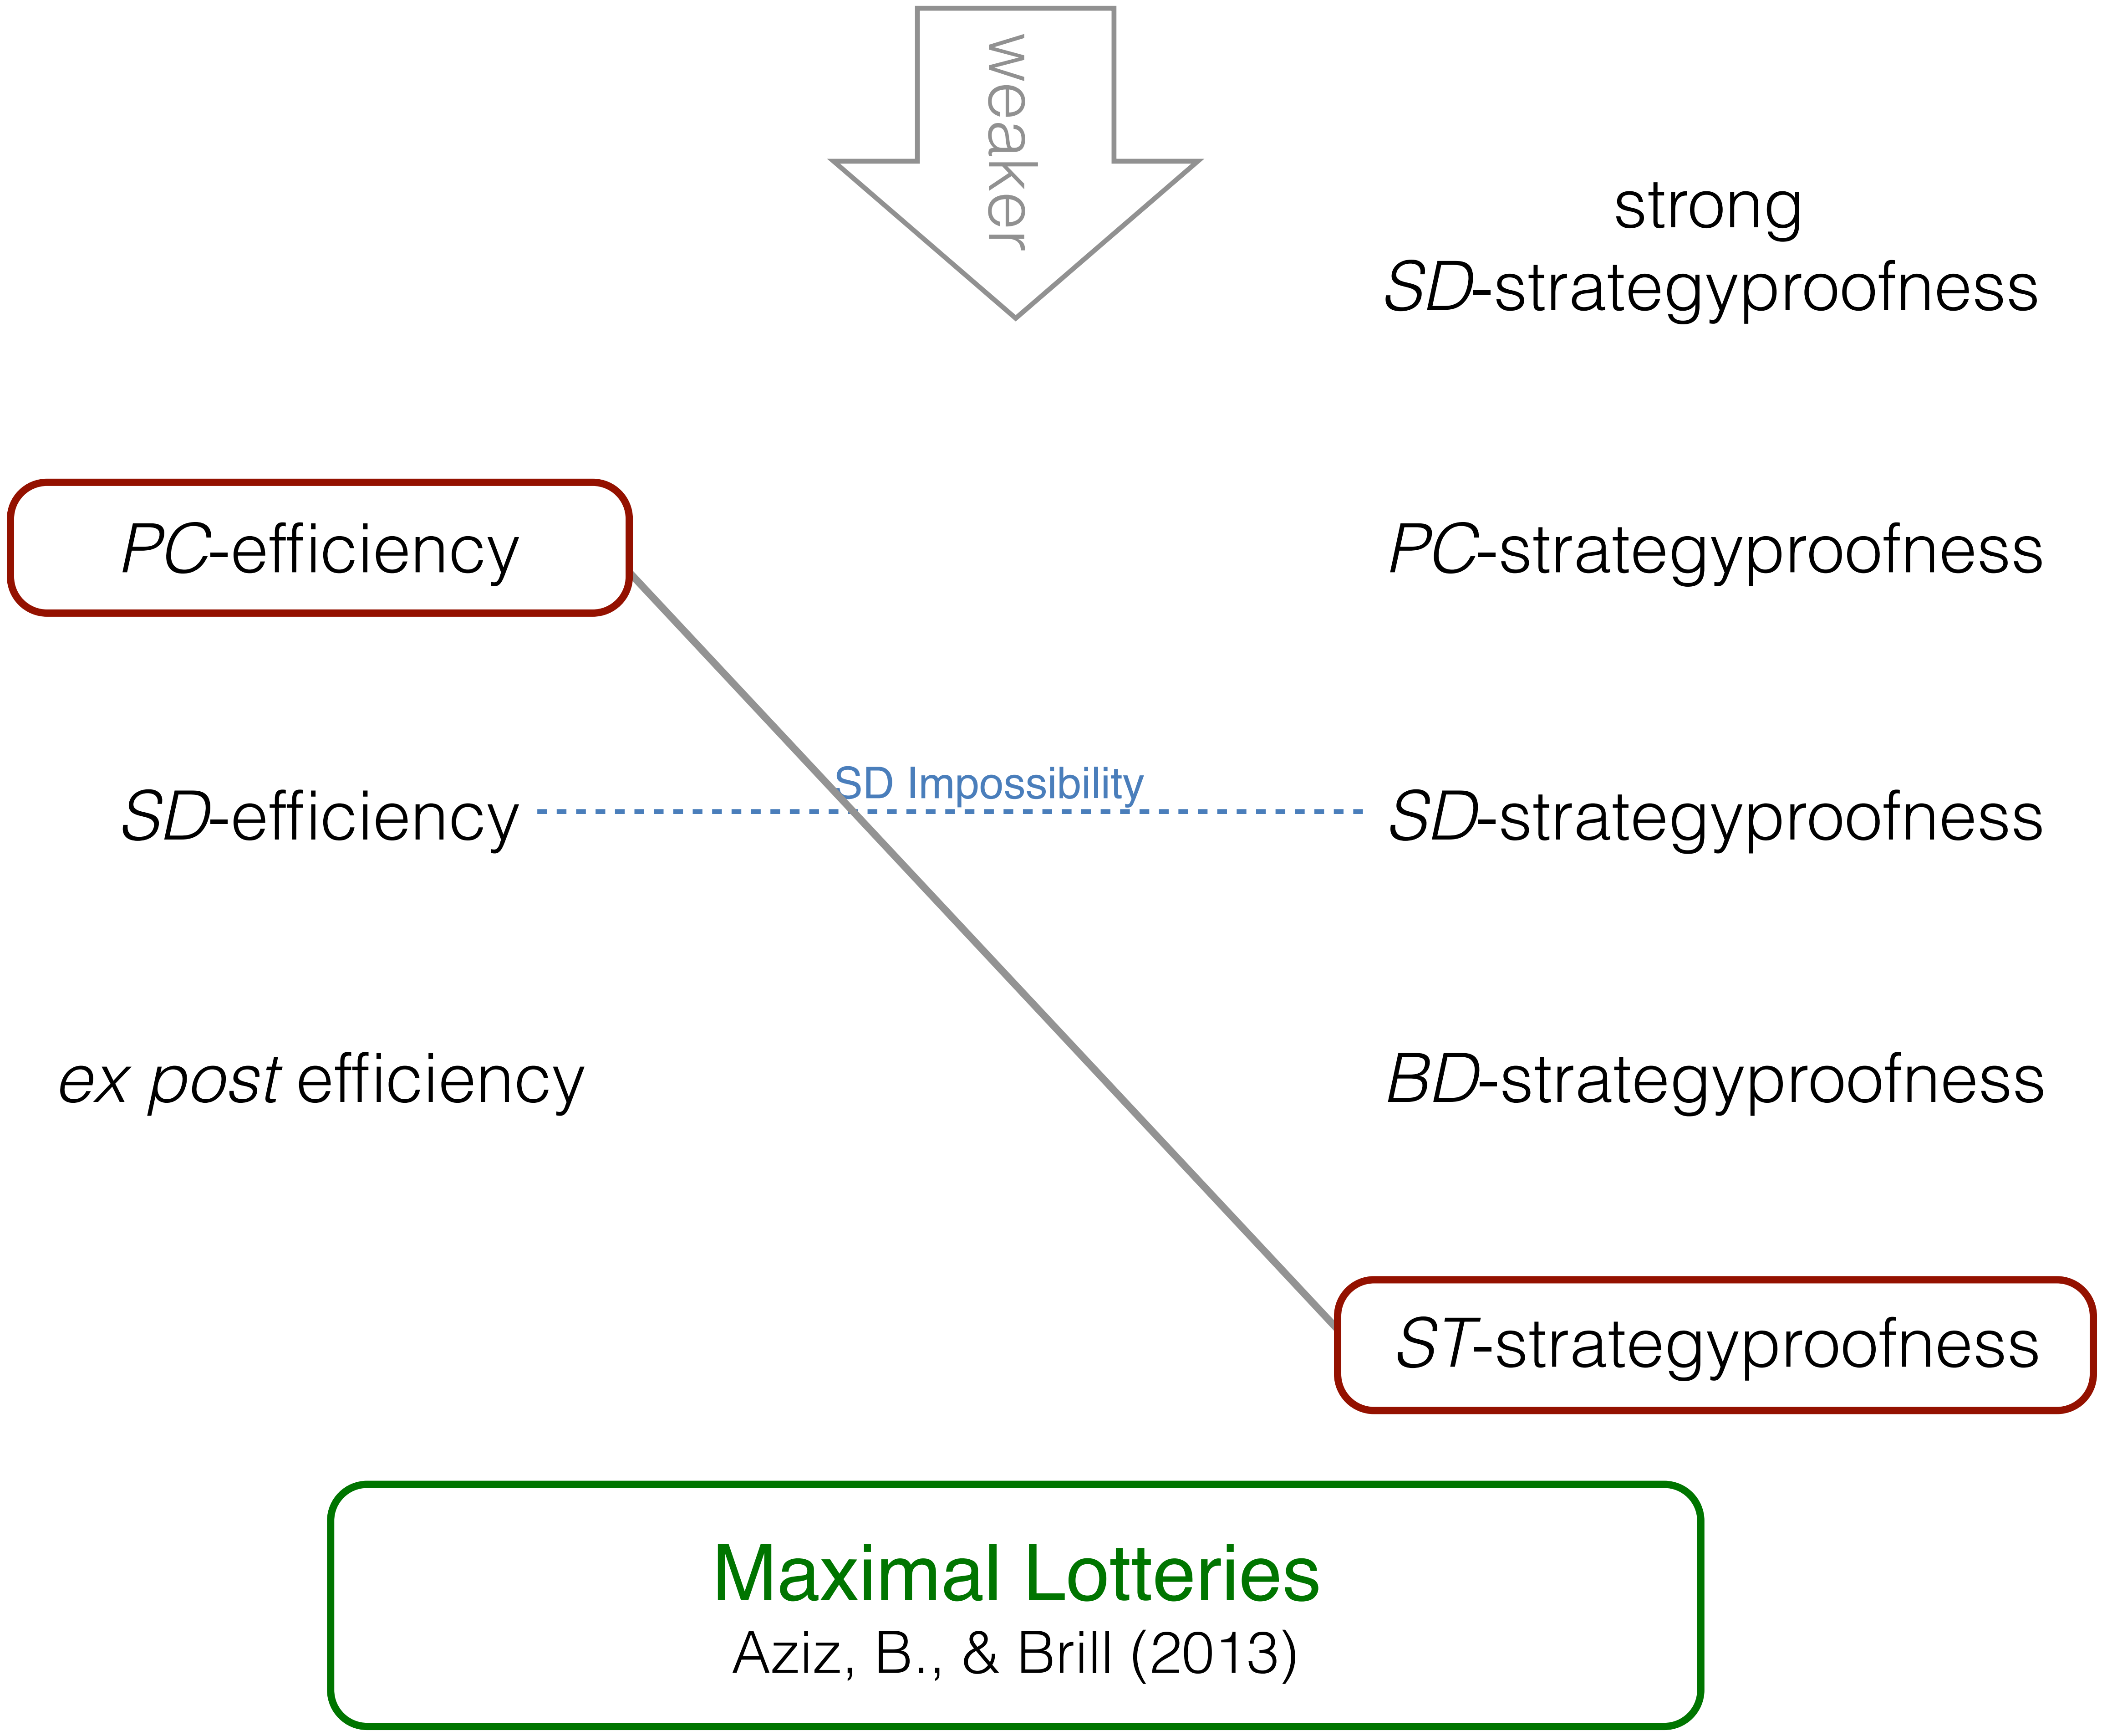
weaker
strong
SD-strategyproofness
PC-efficiency
PC-strategyproofness
SD Impossibility
SD-efficiency
SD-strategyproofness
ex post efficiency
BD-strategyproofness
ST-strategyproofness
Maximal Lotteries
Aziz, B., & Brill (2013)

Conclusion

- No social decision scheme satisfies moderate degrees of efficiency and strategyproofness.
- *RSD* is very strategyproof, but only a little efficient.
- *ML* is very efficient, but only a little strategyproof.
- Further results
  - *RSD* and *ML* are *ST*-group-strategyproof, but not *SD*-group-strategyproof.
  - No anonymous and neutral social decision scheme is *ex post* efficient and *BD*-group-strategyproof, even when preferences are dichotomous.