An impossibility theorem of social choice

http://home.uchicago.edu/~rmyerson/research/schch1.pdf
Can a political institution abolish multiple equilibria?
A variant of Arrow's impossibility theorem says No.

Let N denote a given set of individual voters.
Let Y denote a given set of social-choice options, of which the voters must select one.
We assume that N and Y are both nonempty finite sets.
Let L(Y) denote the set of strict transitive orderings of the alternatives in Y.
Let $L(Y)^N$ denote the set of profiles of preference orderings, one for each voter.
We may denote such a preference profile by a profile of utility functions $u = (u_i)_{i\in N}$, where each u_i is in L(Y). So if the voters' preference profile is u, then the inequality $u_i(x) > u_i(y)$ means that voter i prefers alternative x over alternative y.
($u_i(x) = \#\{y\in T|\ x$ is preferred to y under i's preference in u}.)

A <u>social choice function</u> is any function $F{:}L(Y)^N \rightarrow Y$, where F(u) denotes the alternative in Y to be chosen if the voters' preferences were as in u.
Let $F(L(Y)^N) = \{F(u)|\ u\in L(Y)^N\}$.
Given any game form $H{:}\times_{i\in N} S_i \rightarrow Y$ (where each S_i is a nonempty strategy set for i),
let E(H,u) be the pure Nash equilibrium outcomes of H with preferences u. That is,

$$E(H,u) = \{H(s)|\ s \in \times_{i\in N} S_i, \text{ and, } \forall i\in N, \forall r_i\in S_i, u_i(H(s)) \geq u_i(H(s_{-i},r_i))\}.$$

<u>Theorem (Muller-Satterthwaite)</u> Suppose that a social choice function $F{:}L(Y)^N\rightarrow Y$ and a game form $H{:}\times_{i\in N} S_i\rightarrow Y$ satisfy

$$\#F(L(Y)^N) > 2 \text{ and } E(H,u) = \{F(u)\}\ \forall u\in L(Y)^N.$$

Then there is some h in N such that $u_h(F(u)) = \max_{x\in F(L(Y)^N)} u_h(x)$, $\forall u\in L(Y)^N$.

That is, if an institution H admits more than two possible outcomes and always yields a unique pure-strategy Nash equilibrium, then H must be a dictatorship.
Different democratic institutions may have very different sets of equilibria, but we cannot expect any to abolish multiplicity or randomization of Nash equilibria,
and so democratic outcomes may depend on more than just the voters' preferences.
<u>Lemma (monotonicity)</u> Suppose $E(H,u) = \{F(u)\}\ \forall u\in L(Y)^N$. Then for any u and v,
if $\{(i,y)\in N\times Y|\ v_i(y) > v_i(F(u))\} \subseteq \{(i,y)\in N\times Y|\ u_i(y) > u_i(F(u))\}$, then $F(v) = F(u)$.

<u>Example: the Condorcet cycle.</u> Social options are Y = {a,b,c}, voters are N = {1,2,3}.
$u_1(a)=2 > u_1(b)=1 > u_1(c)=0$; $u_2(b)=2 > u_2(c)=1 > u_2(a)=0$; $u_3(c)=2 > u_3(a)=1 > u_3(b)=0$.
If H is symmetric with respect to social options (<u>neutrality</u>) and voters (<u>anonymity</u>) then its pure-strategy equilibrium outcomes are either E(H,u)=Y (multiple equilibria) or E(H,u)=Ø (only randomized equilibria).

The Probabilistic Voting model: a simple formulation
[For sophisticated probabilistic voting models that also include campaign contributions, see Persson and Tabellini Political Economics (2000) chapters 3 and 5. See also G. Grossman and E. Helpman, "Electoral competition and special interest politics," Review of Economic Studies 63:265-286 (1996), for a model that includes a version of probabilistic voting and campaign contributions.]

Let Y denote the set of social-choice alternatives or policy options for the government.
There are two parties, and each party k in {1,2} can simultaneously choose a policy x_k in Y.
Let us also allow that a party could promise to choose its policy according to any probability distribution σ_k in $\Delta(Y)$.
Each voter has a policy-type i that is independently drawn from a set of types I, getting type i with probability r_i. Each policy y in Y gives some utility $u_i(y)$ to every type-i voter.
In addition, each voter has a net personal bias toward party 1 that is drawn independently from a uniform distribution on the interval $[-\delta,\delta]$. A voter of type i with policy-type i and bias β gets payoff $\beta+u_i(x_1)$ if party 1 wins, but gets payoff $u_i(x_2)$ if party 2 wins.
After the parties have chosen their policy positions x_1 and x_2, each voter votes for the party that offers him the higher payoff, given his policy-type and his bias.
Each party wants to maximize its probability of winning the majority-rule election.

Fact. If both parties choosing the same policy $x_1 = x_2 = x \in Y$ is an equilibrium,
then x maximizes the expected sum of the voters' utility $x \in \text{argmax}_{y\in Y} \sum_{i\in I} r_i\, u_i(y) = Eu_i(y)$.
Proof. When they both choose x for sure, a voter of any policy-type is equally likely to vote for either party, and so each party has an equal probability of winning a majority of the vote.
Now, keeping party 2 at x for sure, suppose that party 1 deviated and promised to choose x with probability $1-\varepsilon$ and some other y with probability ε, given $\varepsilon>0$ and $y\in Y$.
The possibility of changing policy from electing 1 instead of 2 would change type-i's expected utility by the amount $\varepsilon(u_i(y)-u_i(x))$, and so a type-i voter will vote for 1 if his bias β satisfies $\beta+\varepsilon(u_i(y)-u_i(x)) > 0$, that is $\beta > -\varepsilon(u_i(y)-u_i(x))$, which has probability $1/2+\varepsilon(u_i(y)-u_i(x))/(2\delta)$, if ε is small enough so that this formula is between 0 and 1.
So when ε is small, the probability of any randomly-sampled voter voting for party 1 is $1/2 + \varepsilon\sum_{i\in I} r_i\, (u_i(y)-u_i(x))$.
Thus, if $\sum_{i\in I} r_i\, u_i(y) > \sum_{i\in I} r_i\, u_i(x)$ then the ε-probabilistic deviation from x to y would make any randomly sampled voter more likely to vote for party 1 than for party 2, and so (be different voters' votes are independent) the deviating party 1 would get a greater than 1/2 chance of winning the election. But in equilibrium, such deviations from x cannot increase a party's chances of winning, and so we must have $\sum_{i\in I} r_i\, u_i(x) \geq \sum_{i\in I} r_i\, u_i(y)$ for all y in Y.

This result tells us that a convergent pure equilibrium must choose a policy that is a utilitarian optimum, maximizing the expected total utility of all voters.
But this result is somewhat misleading, because such convergent pure equilibria do not generally exist. They exist only when δ is very large, that is, when the effect of policy is small relative to the effect of individuals' biases toward one party or the other.

For example, suppose that $Y=\{a,b,c\}$, $I = \{1,2,3\}$, and $u_i(y)$ is as follows

Type i	$u_i(a)$	$u_i(b)$	$u_i(c)$	r_i
1	2	1	0	0.4
2	0	2	1	0.3
3	1	0	2	0.3
Eu_i:	1.1	1.0	0.9	

So the utilitarian-optimum result is that, if there is a convergent equilibrium where both parties choose the same policy x, it must be the policy x=a, which maximizes voters' expected utility.

But when δ is small, if there are many voters, then policy a almost-surely beats policy b, policy b almost-surely beats policy c, and policy c almost surely beats policy [a], and either party could find a promise that would win with probability greater than 1/2 if it knew what the other party's (possibly probabilistic) promise would be. (Any surely promised randomization in $\Delta(Y)$ could be beaten by another promised randomization that shifts probability from b to a or from c to b or from a to c.)
In the limiting case of $\delta=0$, this case reduces to the Condorcet cycle [ABC cycle] which a unique equilibrium where both parties choose policies randomly in the bipartisan set {a,b,c} as defined by Laffond, Laslier, and Le Breton (1993); see also my Fundamentals of Social Choice Theory survey paper at http://home.uchicago.edu/~rmyerson/research/schch1.pdf
So for small δ, there is no convergent equilibrium where both parties make the same predictable promise. To have a pure convergent equilibrium at policy a here, we must have $\delta > 1.5$.

To see that convergent equilibrium at a requires $\delta>1.5$, consider party 1 deviating to put probability ε on policy c, while party 2 remains at policy a for sure.
60% of the voters prefer c over a, but the 40% type-1s who prefer a care twice as much, and so the fraction of voters whom party 1 gains $0.6(1\varepsilon)/\delta$ is less than the fraction $0.4(2\varepsilon)/\delta$ that party 1 loses by the deviating. But this calculation goes wrong when ε becomes large enough that $(2\varepsilon)/\delta > 1/2$, because then party 1 will have lost all of the type-1 voters, and then further increases in ε can win more voters without losing any more voters. So the equilibrium might be overturned by $\varepsilon=1$ (all probability on c). That is, consider party 1 deviating to c for sure.
The least net pro-1 bias for a type-i voter to support party 1 is then $u_i(a)-u_i(c)$, and so a type-i voter's probability of voting for party 1 is $\max\{0, \min\{1, (\delta - (u_i(a)-u_i(c)))/(2\delta)\}\}$.
So in the whole population, the probability of a voter voting for the deviating party 1 is
$0.4\max\{0,\min\{1,(\delta-2)/(2\delta)\}\} + 0.3\max\{0,\min\{1,(\delta+1)/(2\delta)\}\} + 0.3\max\{0,\min\{1,(\delta +1)/(2\delta)\}\}$
$= 0.4\max\{0,(\delta-2)/(2\delta)\} + 0.3\min\{1,(\delta+1)/(2\delta)\} + 0.3\min\{1,(\delta +1)/(2\delta)\}$.
When $\delta < 1$, this probability of voting for 1 becomes $0+(0.3+0.3)(1) = 0.6 > 1/2$, and so the equilibrium fails.
When $\delta > 2$, this probability of voting for 1 is $1/2-0.4(1/\delta)+(0.3+0.3)(0.5/\delta) = 1/2-0.1/\delta <1/2$, and so the equilibrium does not fail.
When $1\le\delta\le2$, this probability of voting for 1 is $(0.3+0.3)(1/2+1/(2\delta))$ which is $>1/2$ when $\delta<1.5$.
So the equilibrium fails when $\delta<1.5$.

Citizen-Candidate Model T. Besley and S. Coate, "An economic model of representative democracy," Quarterly J of Economics 112:85-114 (1997).

N = {citizens}, Y = {policy space}. For each $i \in N$, $u_i : Y \to \mathbb{R}$ is i's utility for policies.
δ = cost of becoming a candidate. Let θ_i be i's ideal point $\theta_i = \mathrm{argmax}_x\, u_i(x)$.
First, each citizen decides independently whether to become a candidate.
Then all citizens learn $K = \{\text{candidates}\} \subseteq N$, and each votes for one candidate.
The candidate with the most votes is the winner (ties resolved by randomization) and the government policy is the ideal point of the winner.
So if j is winner then each citizen i gets payoff $u_i(\theta_j)$ if $i \notin K$, or $u_i(\theta_j) - \delta$ if $i \in K$.
If $K = \emptyset$ then the outcome is some given x_0 in Y, and i gets $u_i(x_0)$.
The game is analyzed by looking at subgame-perfect equilibria in pure (nonrandom) strategies, after eliminating dominated strategies (voting for the least-preferred candidate) in each subgame after K is determined. The existence of such equilibria can be proven.
We consider cost δ to be small, taking limit as $\delta \to 0$.

An equilibrium in which exactly one candidate enters can only be near or at (as $\delta \to 0$) a Condorcet-winning policy position. Equilibria where exactly two candidates i and j enter (Duverger's equilibria) can exist for any {i,j} such that the number of citizens who prefer i's ideal policy over j's is equal to the number who prefer j's ideal over i's.

Equilibria with three or more tied winners are hard to sustain (for the same reason as in Feddersen AJPS 1992): If a pure-strategy equilibrium generates a tie among k candidates, then no voter can strictly prefer any two of the tied candidates over the k-way randomization, because he could break the tie in favor of whichever candidate he was not expected to vote for (in the eqm).

But we can construct equilibria where three or more candidates enter even though most are expected to lose, because the presence of these spoilers can change the focal equilibrium in the subgame after candidates' entry. Remember, for any pair of candidates, there exists an equilibrium in the plurality-voting election where this pair is considered to be the only serious race, and so everybody votes for the one in this pair whom he prefers.
Consider a simple Hotelling example where Y = [0,100], citizens have ideal points that are distributed uniformly over the interval 0 to 100, and each citizen's policy-payoff is minus the distance of policy from his ideal point.
Pick any x such that $2 < x < 98$. We can construct an equilibrium in which seven candidates enter with ideal points {0,1,2,x,98,99,100}.
On the equilibrium path, the only serious race is 0 versus x, and x wins.
But if any candidate other than x dropped out, then the post-entry subgame equilibrium would switch to one where the only serious race is between two extreme candidates, the least moderate remaining on the side of the unexpected dropout, and the most moderate on the other side.
(E.g.: if 0 dropped out, then the serious race would be between 1 and 98, and 98 would win).
An unexpected extra entrant could be ignored (or could lead to an eqm where the 2 or 98 wins, whichever is worse for the unexpected entrant).

A model of leaders and supporters in contests for power (R,λ,s,c,δ)
The Autocrat's Credibility Problem and Foundations of the constitutional State
http://home.uchicago.edu/~rmyerson/research/foundatn.pdf

An island principality yields income R that can be consumed or allocated by the ruler.
The ruler is the leader who won the most recent battle on the island.
Battles occur whenever a new challenger arrives, at a Poisson rate λ.
(In any time interval ε, P(challenger arrives) = $1-e^{-\lambda\varepsilon} \approx \lambda\varepsilon$ if $\varepsilon \approx 0$.)
A leader needs support from captains to have any chance of winning a battle.
Pr(leader with n captains wins against a rival with m captains) = $p(n|m) = n^s/(n^s+m^s)$.
Let c denote a captain's cost of supporting a leader in battle.
The prince and the captains are assumed to be risk neutral and have discount rate δ.

Consider a leader who has n supporters, but expects all rivals to have m supporters.
(For simplicity, we will always assume stationary expectations about rivals.)
If the leader has promised to give each supporter an income y (as long as the leader rules)
then, when there is no challenger, a supporter's expected discounted payoff is
$U(n,y|m) = (y-\lambda c)/[\delta+\lambda-\lambda p(n|m)]$.
For these captains to rationally give support in battle, we need $p(n|m)U(n,y|m) - c \geq 0$.
The lowest income y satisfying this participation constraint is $Y(n|m) = (\delta+\lambda)c/p(n|m)$
The leader's expected disounted payoff is:
$V(n,y|m) = (R-ny)/[\delta+\lambda-\lambda p(n|m)]$ when he rules with no immediate challenge,
$W(n,y|m) = p(n|m)\, V(n,y|m) = p(n|m)(R-ny)/[\delta+\lambda-\lambda p(n|m)]$ on the eve of battle.

An absolute monarch is one who is released from all constraints of law.
An absolute leader who cheated a supporter would not be punished by anyone else,
although of course the cheated individual might be less likely to support him in the future.
(An absolutist would have no incentive to pay supporters if even those cheated don't react.)
So a leader is absolute when his relationships with all supporters are purely bilateral,
as if supporters have no communication with each other.
Against m, a force of n captains is <u>feasible for an absolute leader</u> iff
there exists some wage rate y such that $y \geq Y(n|m)$ and $V(n,y|m) \geq V(k,y|m)$ $\forall k \in [0,n]$.
First is participation constraint for captains, second is absolutist's moral-hazard constraint.

Let $v(n|m) = V(n,Y(n|m)|m) = [R-nc(\delta+\lambda)/p(n|m)]/[\delta+\lambda-\lambda p(n|m)]$,
and let $w(n|m) = W(n,Y(n|m)|m) = [p(n|m)R-nc(\delta+\lambda)]/[\delta+\lambda-\lambda p(n|m)]$.

<u>Proposition 1.</u> If $n>0$ and y satisfy the feasibility condition for an absolute leader against m, then
there exist $k > n$ such that $v(k|m) > V(n,y|m)$ and $w(k|m) > W(n,y|m)$.
<u>Proof.</u> [Easy if $y>Y(n|m)$.] $Y'(n|m) < 0$. AbsFeas => $V'(n,y|m) \geq 0$. [$'$ = deriv wrt 1st.]
So with $y=Y(n|m)$, $v'(n|m) = V'(n,y|m) - Y'(n|m)n/[\delta+\lambda-\lambda p(n|m)] > 0$.
So an absolute leader could always benefit by commitment to maintain a larger force.

Now suppose captains communicate at court, and a complaint by any captain could switch them to a distrustful equilibrium, where nobody trusts the ruler to reward supporters.
Complaining-only-if-cheated is incentive compatible, as captains expect $U>0$ on eqm path.
With challenges at rate λ and no support, the ruler's expected payoff would be $R/(\delta+\lambda)$.
So we say n is feasible for a leader with a weak court against m iff $v(n|m) \geq R/(\delta+\lambda)$.
$V(0,y|m) = R/(\delta+\lambda)$, so feasible for absolutist => feasible for leader with a weak court.
This court is called "weak" because it cannot change the arrival rate of new challengers.
But when a ruler is known to have no support, immediate challenges may be more likely.
Then loss of confidence at court could lead to a rapid downfall of the leader.
So we say n is feasible for a leader with a strong court against m iff $v(n|m) \geq 0$.

Proposition 2. Suppose that n is feasible for a leader with a weak court against m.
Then $nY(n|m)/R \leq p(n|m)\lambda/(\delta+\lambda)$ and $n \leq R\lambda p(n|m)^2/[c(\delta+\lambda)^2]$.
If $n > 0$ and $s>0.5$ then $m \leq M_0 = [R\lambda(2s-1)^{2-1/s}]/[4s^2c(\delta+\lambda)^2]$.

We may say that a force size m is globally feasible for leaders of some kind (absolute, or with weak courts, or with strong courts) iff m is feasible against m for such leaders.

Proposition 3. Suppose that $s \geq 2/3$.
If n is feasible against m for a weak-court leader and $0 < n \leq m$, then $w'(n|m) > 0$.
So if m is globally feasible for weak-court leaders then $\mathrm{argmax}_{k\geq 0}\, w(k|m) > m$.

We may say that m is a negotiation-proof equilibrium iff $w(m|m) = \max_{n\geq 0} w(n|m)$,
so that any new leader before first battle would want to negotiate the same force size.
By Prop 3, such a negotiation-proof eqm cannot be globally feasible with weak courts.

Proposition 4. When $s\leq 2$, the negotiation-proof equilibrium is $m_1 = Rs/[c(4\delta+2\lambda+s\lambda)]$.
In this eqm, supporters get the fraction $m_1Y(m_1|m_1)/R = 2s(\delta+\lambda)/(4\delta+2\lambda+s\lambda)$ $[\to 1$ as $s\to 2]$.
When $s\geq 0.763$, this equilibrium m_1 is greater than the bound M_0 from Proposition 2, and so an absolutist or a leader with a weak court could not get any support against this eqm.

The courtiers are in a coordination game with multiple equilibria.
With $s>1$, nobody should support a leader who is not supported by anybody else (p too small).
Each wants to support the leader as long as he trusts the leader and enough others are expected to also support the leader. Once a critical mass of supporters has been gathered, before the first battle for power, the leader's speech could make focal the equilibrium with trust among the w-maximizing m supporters.