Affirmative Qualification and Bounded Rejection

Two Questions on One Ballot: A Stagewise Rule for Single-Winner Elections

The Fair Voting System (FVS)

Working Paper, Version 13 — September 2026

Abstract

Electoral institutions pose two different questions. *Which proposals should be under consideration?* asks voters to supply direction. *Among those under consideration, which should be avoided?* asks them to evaluate what has been supplied. Two-round systems pose these in sequence, as a qualification contest followed by an election. This paper develops a rule that poses both on one ballot.

In the Fair Voting System, FVS(0.20, 2), each voter may authorize at most one candidate and actively reject at most two. Affirmative marks construct a qualified set from the share of issued authorizations, $S_j = A_j$ / V > 0.20; active rejections then order the qualified candidates, and the candidate with the lowest rejection rate $R_j = D_j$ / T wins. The signals are never netted. We prove that within a bounded three-state ballot whose unmarked state carries no imputed content, no qualification function satisfying two weak requirements can be a function of the rejection vector alone; the result is conditional on that design class and is not a general impossibility claim.

The paper's principal negative finding concerns the meaning of the negative mark, which an adopting polity must fix in its ballot instructions. Under an *aversion* reading — mark the candidates you most oppose — a candidate entered solely to absorb rejection capacity, holding two per cent of authorizations and incapable of qualifying, can change the winner; Appendix C gives a closed-form instance. Under an *avoidance* reading — mark the candidates you most want to prevent from winning — that manipulation is inert, but the rule reduces to plurality within the qualified set whenever four or fewer candidates qualify and voters authorize inside it (Proposition 8), so the negative channel does no independent work in the configurations that recent large-field elections actually produce. Neither reading dominates. The choice is constitutional, not editorial, and the paper states the dilemma rather than resolving it.

The French presidential election of 2002 is used as a running case. Applying the gate to that year's official first-preference distribution returns no qualifier, symmetrically against an incumbent President, an incumbent Prime Minister and a far-right challenger. The paper reports structural results, counterexamples to majority, Condorcet, Condorcet-loser, participation and reinforcement criteria under both readings of the negative mark, and an explicit account of what the rule does not secure.

Contents

1. Introduction **3**
- 1.1 Ballots as choice and public speech . . . 3
- 1.2 The proposed architecture . . . 3
- 1.3 The normative choice stated plainly . . . 4
- 1.4 Results and limits . . . 4
- 1.5 Plan of the paper . . . 5

2. Two Questions **5**
- 2.1 Political choice as successive compression . . . 5
- 2.2 Two questions, not two kinds of voter . . . 5
- 2.3 Proposal power and the reversion outcome . . . 6
- 2.4 Three institutional acts . . . 7
- 2.5 Same-type aggregation, cross-type sequencing . . . 7
- 2.6 Conditional delegation and the consequence of silence . . . 7
- 2.7 The ordering theorem . . . 7
- 2.8 Affirmative screening relocates the ignorance default . . . 8

3. A Running Case: France, 2002 **9**
- 3.1 The first round . . . 9
- 3.2 Comparison across five elections . . . 10

3.3 A quota gate and a rank cutoff diverge . . . 10
3.4 The second round and the recording problem . . . 10
3.5 One electorate, two questions . . . 11
3.6 What the case does not show . . . 11

4. Relation to Existing Rules and Institutions **11**
4.1 Rankings, approval, and public meaning . . . 11
4.2 Two-round systems . . . 11
4.3 The constructive vote of no confidence . . . 12
4.4 Negative ballots and threshold aggregation . . . 12
4.5 The architectural comparison . . . 13
4.6 Functional separation rather than motivational purity . . . 13

5. Formal Definition of FVS **13**
5.1 Domain and ballots . . . 13
5.2 Counts and denominators . . . 14
5.3 Qualification and selection . . . 14
5.4 The complete public output . . . 14
5.5 The required definition of the negative mark . . . 14
5.6 No validity threshold on V/T . . . 15

6. Structural Properties **15**
6.1 Affirmative normalization and the proposal cap . . . 15
6.2 Threshold–budget compatibility . . . 16
6.3 Directional abstention neutrality . . . 16
6.4 Stagewise non-compensation . . . 16
6.5 The aggregate rejection bound . . . 16
6.6 Own-signal monotonicity . . . 16
6.7 Bounded rejection, and what each reading of the negative mark implies . . . 17
6.8 Ballot-space complexity . . . 17
6.9 Entry, positive-pool externality, and reinforcement . . . 18

7. Strategic and Axiomatic Stress Tests **18**
7.1 No privileged definition of sincerity . . . 18
7.2 Two analytical mappings, and why the results diverge . . . 18
7.3 Majority failure (aversion mapping) . . . 19
7.4 Condorcet-loser failure (both mappings) . . . 19
7.5 Participation failure (aversion mapping) . . . 19
7.6 A profitable positive reallocation (aversion mapping) . . . 19
7.7 Reinforcement and field dependence . . . 20

8. Information, Viability, and the Public Record **20**
8.1 Qualification is not an awareness test . . . 20
8.2 The viability-visibility condition . . . 20
8.3 Near-threshold speech is not erased . . . 20
8.4 What the winner's mandate means . . . 20
8.5 Polarity fidelity, not complete semantic fidelity . . . 21

9. Normative Trade-offs and Institutional Limits **21**
9.1 Constructive minorities and negative majorities . . . 21
9.2 Positive abstention is free but not costless . . . 21
9.3 Fragmentation and two directions of proliferation . . . 21
9.4 Negative marks are comparative, not vetoes . . . 22
9.5 Candidate-set rejection: three constraints and an analytical baseline . . . 22
9.6 Scope of the parameterization . . . 22
9.7 No mechanical claim about negative campaigning . . . 23
9.8 Upstream inequality remains . . . 23

10. Research Agenda **23**
10.1 The frequency of qualified-set sizes . . . 23

10.2 The semantics of the negative mark . . . 23
10.3 Recognition, independently measured . . . 23
10.4 V/T as an open measurement . . . 23
10.5 The qualification gate and parameter sensitivity . . . 23
10.6 The entry game . . . 24
10.7 Continuation rules . . . 24
10.8 Campaign effects . . . 24
10.9 Public-record value and pilots . . . 24

11. Conclusion **24**

References **25**

Appendix A. Compact Executable Specification **26**

Appendix B. Reporting Template **27**

Appendix C. A Closed-Form Instance of Absorbing Entry **27**
C.1 Setup . . . 27
C.2 Configuration 1: no entrant . . . 28
C.3 Configuration 2: entrant X, aversion reading . . . 28
C.4 Configuration 3: entrant X, avoidance reading . . . 28
C.5 What the example establishes and what it does not . . . 29

Keywords: voting systems; ballot design; agenda control; proposal power; approval and disapproval voting; negative voting; threshold aggregation; two-round systems; bounded expression; public choice

JEL: D71, D72

1. Introduction

1.1 Ballots as choice and public speech

Democratic elections must end in an executable decision, but the ballot producing that decision is also a public communication device. How a ballot asks citizens to act affects both who wins and what the resulting tally can truthfully be said to mean.

Every ballot compresses. Voters hold judgments about policies, competence, character, coalitions, risks and priorities; an election reduces these to a small number of marks and then to one officeholder. The design question is not whether information is lost. It is which distinctions the institution preserves, which it merges, and in what order the surviving distinctions enter the collective decision.

Consider a voter who regards candidate A as bad and candidate B as worse. In a positive-only contest, the competitively effective candidate-directed action against B may be a vote for A. This is not a claim that voting for A is a dominated or insincere act; the voter may genuinely authorize A as the best executable option available. The recording problem is narrower: the ballot has no separate candidate-directed action with which to state opposition to B. The tally records +A and cannot also record the anti-B act that helped motivate it.

Section 3 supplies an unusually clean empirical instance. In the second round of the 2002 French presidential election the winner received 82.21% of votes cast, having received 19.88% in the first round three weeks earlier. The intervening 62 percentage points contained very little authorization of the winner.

A direct negative channel expands the message alphabet. It makes it possible to record −B without mechanically adding +A. This does not purify motives. The institutional gain is not sincerity certification but typed expression: a positive mark, a negative mark and an omission are not forced into a single scalar sign.

1.2 The proposed architecture

The Fair Voting System is a one-stage, single-winner rule for contests with at least four formally nominated candidates. In FVS(0.20, 2), every allocation is optional. A voter may issue one affirmative authorization or none; as many as two active candidate-specific rejections or fewer; and leave every other candidate unmarked. The marks are submitted simultaneously.

Aggregation is stagewise. Positive marks construct a qualified proposal set; negative marks then order its members. The two are strategically connected, because both can affect the result, but they are not arithmetically compensable. Once two candidates qualify, additional authorization cannot offset a strictly higher rejection total; affirmative share breaks only an exact rejection tie.

The contribution claimed is not the invention of negative voting, best-worst elicitation, or non-compensatory approval–disapproval aggregation. Those have established literatures, and Section 4.4 identifies a characterized threshold rule whose selection logic is close to the one used here. What is claimed is an argument from the distinction between two institutional questions to a specific assignment of channels, a theorem establishing that assignment within a restricted design class, a specification implementing it under scarce budgets, and — the paper's sharpest result — a demonstration that the two admissible readings of a scarce negative mark generate a genuine dilemma rather than a choice of wording.

1.3 The normative choice stated plainly

FVS treats a coalition of opposition differently from a coalition capable of governing. A group may agree that candidate A is unacceptable while having no common alternative that any sufficient share is willing to authorize. That group has expressed politically important opposition, but it has not supplied an executable direction. FVS records the opposition; it does not convert it into a positive proposal and does not allow it, by itself, to manufacture a qualified rival.

This yields a deliberately demanding limiting case. Suppose 100 participants take part, one authorizes A, and the other 99 issue no authorization but all reject A. Then V = 1, $S_A = 1$, and Q = {A}. Candidate A wins despite a rejection rate of 99%. The outcome does not mean that A has majority support. It means that A is the only affirmatively qualified direction and that the 99-person opposition did not form a qualified alternative.

> **An opposition majority is not an action majority. To replace the only qualified direction, opponents must coordinate around at least one qualified alternative.**

Section 4.3 gives the established institutional analogue, the constructive vote of no confidence, and shows that FVS demands strictly less.

Two qualifications belong with the principle.

First, the demand is **a demand for formal concentration, and the rule supplies no technology for meeting it.** Section 9.3 shows that a tendency commanding 54% of issued authorizations, divided equally among three candidates at 18% each, qualifies nobody while a rival at 46% qualifies. Identifying three candidacies as "one direction" would require either an administrator empowered to make that judgment, which the design rejects throughout, or a different ballot permitting set-valued or ordinal support. Coordination into a qualifying candidacy is work for candidates, parties and voters, performed against a threshold public and fixed in advance.

Second, **the position of the line is a convention, and the paper does not claim otherwise.** Proposition 1 derives α from an upper-bound requirement: $0.20 \le \alpha < 0.25$ is precisely the interval guaranteeing a cap of four while permitting four to qualify. That argument concerns how many proposals may stand. It does not establish that 19.58% of issued authorizations differs in political substance from 20.01%; in the 2017 French first round those figures separated the fourth and third candidates.

1.4 Results and limits

Theorem 1 (Section 2.7). Within the class of typed, non-imputing ballot rules, no qualification function satisfying *no standing without authorization* and *affirmative unanimity* can be a function of the rejection vector alone. **The result is conditional on that class**; Section 2.7.3 identifies designs outside it for which a rejection-first architecture remains coherent.

Proposition 8 (Section 6.7) and Appendix C. The two admissible readings of the negative mark are not interchangeable. Under the aversion reading, an entrant who cannot qualify can change the winner by absorbing rejection capacity. Under the avoidance reading, that manipulation is inert, but if every participant authorizes a qualified candidate and $q \le d + 1$, then $D_j = T - A_j$ for every qualified j and the rule selects the plurality winner within Q. For d = 2 this covers $q \le 3$, and Section 3.2 reports that in five French presidential first rounds the qualified set never exceeded three.

Structural results in Section 6: affirmative shares sum to one, so a strict 20% threshold admits at most four candidates; $\alpha \geq 1/(d+3)$ admits at most d + 2; positive abstention is neutral with respect to the distribution of positive marks already issued; aggregation is non-compensatory; $\sum_j R_j \leq d$; the rule is mechanically monotone in a winner's own positive and negative signals.

Four limits are stated at the outset.

The gate does not certify awareness. Section 8.1 gives a case in which FVS rewards obscurity despite the gate. Section 2.8 states the general form: affirmative screening does not eliminate error arising from unfamiliarity, it relocates the default from automatic clearance to non-qualification.

The gate compels exposure only up to the threshold. Once a candidate qualifies, affirmative share affects the outcome only through an exact-tie break, while broader exposure raises the probability of entering some voter's rejection priorities. Within the counting rule the incentive to seek wider authorization saturates near the threshold.

The scope condition concerns concentration, not the number of candidates. FVS(0.20, 2) requires an electorate in which affirmative authorization can concentrate above 20%. A simulation drawing first-place shares uniformly reports near-certain failure for any m ≥ 6, which is an arithmetic consequence of the uniformity assumption. Section 3.2 gives five real distributions with ten to sixteen candidates, four of which yield a non-empty qualified set. **Five observations from one country and one nomination system are an illustration and not a distribution**; Section 10.1 states the cross-national tabulation required to convert this scope claim into evidence.

No acceptable range for V/T is specified, and this is a position rather than an omission. Section 5.6 explains: any such range would re-read positive abstention as collective negation, and Section 2.6 declines to make that inference at the individual level. V/T is measured and published inseparably from the result; it is not a validity condition.

1.5 Plan of the paper

Section 2 distinguishes the two questions an electoral institution can pose, derives the assignment of channels and the bounded negative budget, and proves the ordering theorem. Section 3 presents the 2002 French election as a running case. Section 4 locates FVS in the literature, including two-round systems, the constructive vote of no confidence, and threshold aggregation. Section 5 gives the formal rule and the required definition of the negative mark. Section 6 derives the structural results, including the collapse proposition. Section 7 reports strategic and axiomatic failures under both readings. Section 8 examines information and the public record. Section 9 states normative and institutional trade-offs. Section 10 gives a research agenda. Section 11 concludes. Appendix A gives an executable specification, Appendix B a reporting template, and Appendix C the closed-form entry example.

2. Two Questions

2.1 Political choice as successive compression

Let Ω denote the politically conceivable set of proposals. Parties, nomination rules, finance, organization, media access and legal requirements compress Ω into the formally supplied candidate set C. An electoral rule compresses C into a qualified set Q, and finally into a winner w:

$$\Omega \to C \to Q \to w.$$

FVS operates only on the last two mappings. It does not equalize the upstream conversion of ideas into candidacies.

A candidate is treated here as an executable electoral proposal: a person, platform, expected coalition and governing package presented for authorization. This does not imply that voters know the package completely. It identifies the unit on which the ballot acts.

2.2 Two questions, not two kinds of voter

The organizing distinction of this paper concerns what an institution asks, not what voters are like.

The supply question. *Which proposals should be under consideration?* Answering it requires nominating a direction, by an act that adds standing to something.

The evaluation question. *Among the proposals under consideration, which should be avoided?* Answering it requires assessing supplied options against what happens otherwise, by an act that withholds or opposes.

These are different questions, requiring different information and different operations. An institution can pose either, both, or neither, and can pose them in sequence or together.

Nothing in this distinction concerns a fixed disposition of voters. The same electorate answers whichever question the institution poses. The relevant variable is the function of the electoral event, and Section 3.5 documents one electorate answering both questions within three weeks.

2.2.1 Enlarging the menu does not change the question

Increasing |C| does not convert an evaluation question into a supply question. Sixteen candidates on a ballot asking "which one do you want" still asks the voter to nominate a direction; sixteen candidates on a ballot asking "which do you most want to prevent" still asks for an evaluation. The question is set by the rule, not by the size of the menu. A rule cannot obtain answers to the supply question merely by lengthening the list.

2.2.2 Avoidance across several options is a prioritization problem

With one supplied option, avoidance is a binary evaluation. With several, a voter wishing to avoid more than one must decide which avoidance matters more.

This yields the argument for bounding the negative budget, and it depends on no behavioural assumption:

> **Avoidance across multiple options requires prioritization, and prioritization carries information only when the budget is scarce. With unlimited rejections, "what I am trying to avoid" collapses into "whom I dislike", and the ordering content of the negative channel vanishes.**

Scarcity is therefore constitutive of what the negative channel records.

A cost accompanies this benefit and should be recorded here rather than discovered later. A bounded budget makes the positions within a voter's rejection priority set a rival good: an entrant who occupies one displaces someone from it. That is the mechanism of the manipulation analyzed in Section 9.3.2 and Appendix C, and it exists precisely because the budget is bounded. Whether the displacement is exploitable turns on the definition adopted under Section 5.5. Scarcity is therefore both the source of the channel's information and the source of its vulnerability.

2.2.3 Both questions are answerable, and voters answer the supply question when asked

An earlier version of this project suggested that the affirmative channel asks something voters are structurally disinclined to supply. The 2002 French first round is evidence against that suggestion.

In that round, sixteen candidates received first preferences and roughly a quarter of all first preferences went to candidates polling below 5% — candidates who could not reach the second round and whose supporters had ample public information to know it. Those voters were answering the supply question, at scale, in full knowledge that their answer would not determine the officeholder. **When an electoral event is framed as a qualification contest, voters supply direction.**

This bears on V/T, the share of active participants issuing an authorization. There is no basis in this evidence for expecting the affirmative channel to go systematically unused. What the evidence does not show is what happens when declining to authorize becomes a distinct available option, as it is under FVS but was not in 2002. V/T therefore remains a quantity to be measured rather than assumed.

2.3 Proposal power and the reversion outcome

The evaluation question has a well-studied formal setting. In the controlled-agenda model of Romer and Rosenthal (1978), an agenda setter proposes a single alternative, voters accept or reject, and rejection produces a fixed reversion outcome. The setter's power does not come from any inability of voters to refuse. It comes from the position of the reversion: voters who unanimously regard the proposal as far from their ideal accept it whenever the reversion is worse. The same structure appears in the pirate division problem, where backward induction yields a division overwhelmingly favourable to the proposer because responders evaluate rationally against an unfavourable continuation.

Two consequences carry into the design. **Unanimous opposition to a proposal is not equivalent to the capacity to produce an alternative**; that capacity is proposal power, distributed in the $\Omega \to C$ mapping rather than by the counting rule. And **the reversion outcome is a design variable of the first order**: when FVS returns candidate-set rejection, what follows is a reversion whose position determines how much proposal power the rule leaves with incumbent agenda setters. Section 9.5 treats the choice.

A bridge is needed here, and the paper states it rather than leaving the inference implicit. Both models concern a single proposal compared against a reversion. An election's candidate set is already the product of many proposals from many actors, with no single agenda setter. What these models establish is therefore narrow: **they show why opposition is not equivalent to direction.** They do not show how qualification should be constructed for a multi-candidate field. That question is addressed by Theorem 1, within a design class whose boundaries Section 2.7.3 makes explicit, and by the default-relocation argument of Section 2.8. The controlled-agenda literature motivates the problem; it does not settle the construction.

2.4 Three institutional acts

Affirmative authorization. A positive mark means: *this is a direction I am willing to authorize as a possible collective result*. It need not identify the voter's private favourite or enthusiasm.

Active rejection. A negative mark means: *this candidate is one of the limited targets against which I choose to register active opposition*. Because the budget is bounded, the mark identifies priority, not an absolute veto or a complete bottom set. Section 5.5 requires the adopting polity to fix which of two readings this priority carries.

Omission. An unmarked candidate receives no candidate-directed signal. Omission does not establish approval, disapproval, indifference, awareness, or psychological acceptability.

2.5 Same-type aggregation, cross-type sequencing

Positive signals aggregate with positive; negative with negative; and the two are not subtracted from one another, though joined sequentially. This is not independence — the positive stage determines the comparison domain — and it is not a signed score. Ten authorizations and nine rejections do not become a net value of one, because "10 for, 9 against" communicates polarization that a net total suppresses.

Direction and constraint can be contributed by different participants. A group deciding where to eat may receive one destination proposal and, from others, restrictions no destination should violate. **The analogy must be bounded immediately.** A dietary restriction is binding: it removes options from the feasible set. An FVS rejection does not. It does not remove a candidate from Q, does not modify a platform, and cannot stop a sole qualifier. It competes; it does not constrain.

2.6 Conditional delegation and the consequence of silence

A participant who casts negative marks but no positive mark declines to contribute to affirmative agenda formation. Under FVS this is conditional delegation: *others may supply the direction; I retain the ability to prioritize opposition within my negative budget*. If a candidate is left unmarked, the mechanism does not infer psychological acceptance.

No voter is forced to say that a candidate is good. But the collective rule is not required to count the absence of a positive proposal as a vote against every proposal offered by others, which would convert positive silence into phantom opposition and dilute the scarce authorizations actually issued. **This refusal to read silence as negation operates at every level of aggregation**; Section 5.6 draws the consequence for V/T.

2.7 The ordering theorem

2.7.1 The design class

Consider a one-stage, optional, candidate-directed ballot on a frozen candidate set C with participant set of size T. Each ballot assigns to each candidate exactly one of three observable states: affirmative, rejected, or unmarked. Write A = (A_j) for affirmative counts and D = (D_j) for rejection counts. A **qualification function** is a map $Q : (A, D) \mapsto 2^C$.

C1 (Typed action). Affirmative authorization and active rejection are distinct public acts. Q takes A and D as separate arguments and does not reduce them to a single net scalar.

C2 (Non-imputation). The unmarked state is not translated into either act. Q depends on the profile only through (A, D).

Axiom N (No standing without authorization). If no participant issues an affirmative mark, Q(A, D) = ∅ for every D.

Axiom U (Affirmative unanimity). If at least one affirmative mark is issued and every issued affirmative mark names candidate a, then a ∈ Q(A, D).

Neither axiom presupposes that qualification is affirmative.

2.7.2 The result

> **Theorem 1.** Let Q be a qualification function on a design class satisfying C1 and C2, with |C| ≥ 2 and at least one participant. If Q satisfies Axiom N and Axiom U, then Q is not a function of D alone.

Proof. Suppose Q(A, D) = g(D) for all realizable (A, D). Fix a ∈ C and construct a rejection vector D^* with $D^*_a = 0$, realizable by participants rejecting only candidates other than a; such a vector exists whenever |C| ≥ 2.

Profile 1. Every participant issues the rejections in D^* and no affirmative mark. Then A = 0 and the rejection vector is D^*. By Axiom N, Q = ∅, so g(D^*) = ∅.

Profile 2. Every participant issues the same rejections and additionally authorizes a. This is legal: the ballot forbids authorizing and rejecting the same candidate, and $D^*_a = 0$ by construction. The rejection vector is again D^*, and every affirmative mark names a. By Axiom U, a ∈ Q, so a ∈ g(D^*).

Then g(D^*) = ∅ and a ∈ g(D^*), a contradiction. ∎

> **Corollary 1.** Under C1, C2, N and U, the qualified set must depend on the affirmative vector. Active rejection can operate only as a selection device within a set whose membership the affirmative channel has helped determine.

The result is elementary. It is stated as a theorem because the conclusion it licenses has often been asserted without one. It establishes an ordering, not a particular gate.

2.7.3 What Theorem 1 does not establish

The antecedents are substantive. If a ballot admits a fourth recorded state — *unknown* or *not evaluated* — C2 no longer forces omission and ignorance to share an observable, and a rejection-first architecture may be coherent. If a constitution treats the status quo as an executable proposal with standing, Axiom N loses its force. If a negative supermajority may trigger a new nomination round, the rule is not a qualification function in this sense.

These designs are not shown impossible; they lie outside the class. **Theorem 1 is not a general impossibility result for negative qualification**, and no statement in this paper should be read as one. It says nothing about whether the affirmative gate should be a quota, a rank cutoff or any other form, and nothing about the parameter value.

2.8 Affirmative screening relocates the ignorance default

A voter who has not encountered a candidate withholds a rejection under one architecture and withholds an authorization under the other. Neither ballot recovers the missing evaluation. Alcantud and Laruelle (2014) make the underlying observation in characterizing dis&approval voting: voters are often indifferent or insufficiently informed about some candidates, so the set a voter has not crossed out contains both those she endorses and those about whom she has no view.

Screening channel	A participant unfamiliar with j	Default consequence for j	Error direction
Rejection screens	issues no rejection	counted toward clearance; j advances	unfamiliar candidates advance

Screening channel	A participant unfamiliar with j	Default consequence for j	Error direction
Affirmative screens	issues no authorization	no qualification credit; j does not advance	unfamiliar candidates fail to advance

Choosing the affirmative screen is a choice about which error the institution bears. Three considerations support it.

Remediability. A candidate failing to qualify because too few participants encountered her can seek exposure and stand again; the remedy lies partly within her control and entirely before the decision is executed. A candidate advancing because too few encountered her may take office, and no remedy operates after the fact.

Legibility of the failure. Under an affirmative screen the failure appears in the record: S_j is low, and low S_j is published. Under a rejection screen it is invisible — a low rejection count produced by obscurity and one produced by broad acceptability are the same number.

What qualification commits the electorate to. Admitting an unexamined option to Q places it in the set from which an executable decision will be drawn.

The cost must be recorded plainly. An affirmative screen systematically disadvantages new entrants, low-resource candidates, and proposals whose support is real but not yet widely encountered. This is a distributive consequence of the design, and Section 9.8 treats it as such.

3. A Running Case: France, 2002

A qualification applies to every computation in this section and is repeated wherever the results are cited. FVS ballots were never cast in 2002; the authorizations and rejections the rule would elicit do not exist. What follows applies the FVS gate to the official first-preference distribution, which coincides with S_j under the canonical mapping of Section 7.2. The exercise shows where the gate falls on a real affirmative distribution of that shape. It does not reconstruct how voters would have allocated authorizations had a rejection channel been available, and — given Sections 5.5 and 6.7 — it does not predict which of the two admissible readings of the negative mark would have produced which winner.

All vote shares in this section are drawn from the published results of the French Ministère de l'Intérieur. Figures should be re-verified against the official proclamation of the Conseil constitutionnel, with the retrieval date recorded, before circulation.

3.1 The first round

Sixteen candidates stood. Leading first-preference shares:

Candidate	Share
Chirac	19.88%
Le Pen	16.86%
Jospin	16.18%
Bayrou	6.84%
Laguiller	5.72%
Chevènement	5.33%
Mamère	5.25%
Nine further candidates, each below 5%	*≈ 24% combined*

Applying the gate returns **Q = ∅**.

The salient feature is the dispersion, not the margin. Roughly a quarter of first preferences went to candidates below 5%, and none of the three leading candidates reached 20%. The verdict "no proposal reached the required concentration of authorization" is a description of that dispersion. It does not rest on Chirac's shortfall of 0.12 percentage points, and the paper does not claim that 19.88% differs in political substance from 20.01%. Had the

threshold been placed at 15%, the gate would have admitted three candidates rather than none, and the character of the round — an electorate whose affirmative direction was scattered across sixteen options — would be unchanged.

The verdict is symmetric. The rule does not filter extremism and makes no judgment about any platform. It applies one public line to an incumbent President, an incumbent Prime Minister and a far-right challenger alike, and all three fail it. A rule screening for extremism would immediately face the question of who defines extremism and by what authority.

3.2 Comparison across five elections

Applying the same computation, under the same qualification, to five French presidential first rounds:

Year	Candidates	Highest first-preference share	Resulting $\|Q\|$
2002	16	19.88%	**0**
2007	12	31.18%	2
2012	10	28.63%	2
2017	11	24.01%	3
2022	12	27.85%	3

Three observations follow, the third of which is load-bearing for Sections 6.7 and 7.

The gate operates at realistic field sizes. With ten to sixteen candidates, four of five distributions yield a non-empty qualified set. The binding requirement is concentration, not the number of candidates. **Five observations from one country, one nomination system and one first-round rule are an illustration, not a distribution**; Section 10.1 sets out the cross-national tabulation required before this is treated as evidence.

Candidate-set rejection is a live output, not an edge case.

The qualified set was never of size four. Here $|Q| \in \{0, 2, 2, 3, 3\}$. Proposition 8 shows that under the avoidance reading of the negative mark the rule reduces to plurality within Q whenever $q \leq 3$, and Section 6.7 shows that the structural rationale for $d = 2$ binds only at $q = 4$. If $q \leq 3$ is the typical configuration, both facts apply to the typical case.

3.3 A quota gate and a rank cutoff diverge

The first round of a French presidential election is a qualification contest with a **rank cutoff**: the top two advance regardless of their shares. FVS uses a **quota**. In 2002 the difference is the whole story.

A rank cutoff necessarily produces exactly two qualified candidates. In 2002 it produced a runoff between two candidates whose combined first preferences were approximately 36.7%. **That pairing was generated by the rule, not selected by the electorate**, and the second-round figure of 82.21% is a downstream consequence of a comparison set the electorate had not endorsed.

A quota gate does not have this property. It produces however many candidates clear the line, including none. Whether that is preferable depends on what the institution should do when the electorate has not concentrated its authorization, which is the question of Section 9.5. The two devices are not interchangeable, and the comparison does not depend on where the quota is set.

3.4 The second round and the recording problem

In the second round Chirac received 82.21% and Le Pen 17.79%. Chirac's first-round share had been 19.88%.

The gap of roughly 62 percentage points is the recording problem in its clearest available form. It did not represent authorization acquired in three weeks; it represented opposition to the alternative, expressed through the only candidate-directed action the ballot supplied. Contemporary campaigning was explicit about this, and no informed observer treated the figure as a measure of support for the winner.

That last fact invites an objection worth answering directly. If everyone understood that 82.21% was not an endorsement, what harm did the recording problem do?

Two answers. First, informal understanding is not an institutional record. The official result is 82.21%; it enters archives, citations and claims of governing authority. Second, and more important: **informal correction works only when the distortion is conspicuous.** It was conspicuous in 2002 because the alternative was Le Pen. Every lesser-evil vote in every ordinary election does the same thing without a conspicuous alternative to make it visible. 2002 is the rare case in which a routine feature of positive-only ballots became impossible to ignore.

3.5 One electorate, two questions

The first round was a qualification contest, and voters answered the supply question: they distributed authorization across sixteen candidates, roughly a quarter of it to candidates who could not advance. The second round was an election between two supplied options, and voters answered the evaluation question. The same electorate answered both within three weeks. **What changed was the question the institution posed.**

A two-round system poses the questions in sequence, at the cost of two mobilizations and a first round in which voters must guess how their answer will be used. A one-stage rule cannot separate the questions in time. **FVS separates them by channel instead:** authorization answers the supply question, bounded rejection answers the evaluation question, and both are answered on one ballot. Section 4.2 states the corresponding cost.

3.6 What the case does not show

FVS would not have prevented the dispersion. The nine candidates below 5% would still have stood. What the rule declines to do is package a dispersed affirmative distribution as an 82% mandate. The claim is about the record, not about the outcome.

FVS would not have produced a President. France required a head of state in May 2002. The rule returns candidate-set rejection. Section 9.5 supplies an analytical baseline continuation and works the case through under it.

The 82.21% figure did not deceive contemporaries. The argument rests on the difference between an institutional record and an informal understanding.

The computation is not a reconstruction of FVS ballots. See the qualification opening this section.

4. Relation to Existing Rules and Institutions

4.1 Rankings, approval, and public meaning

Arrow's impossibility theorem concerns aggregation of individual orderings into a social ordering under specified axioms (Arrow, 1951). FVS neither resolves nor evades it; it changes the elicitation object and produces a single winner.

A ranking $A \succ B \succ C \succ D$ locates candidates relative to one another but does not by itself establish an absolute acceptability boundary. Kruger and Sanver (2021) study difficulties in combining ordinal and evaluative information within a single rule. Approval voting allows a voter to approve any subset (Brams & Fishburn, 1978; Alós-Ferrer, 2006); non-approval can combine rejection, neutrality, strategic truncation and omission. Work on voting as statement-making recognizes expressive value apart from narrow pivotality (Aragones, Gilboa, & Weiss, 2011). Majority Judgment (Balinski & Laraki, 2011) is non-compensatory in a different sense, aggregating common-language grades by median.

4.2 Two-round systems

The closest institutional relative is the two-round runoff, and the comparison should be explicit because FVS poses in one event the two questions a runoff poses in two.

	Two-round runoff	FVS
Qualification	First round; rank cutoff (usually top two)	Quota on issued authorizations
Size of qualified set	Fixed by the rule	Varies; may be 1, may be empty
Selection	Second round, **with Q known**	Same ballot, **with Q unknown**

	Two-round runoff	FVS
Cost	Two mobilizations, two elections	Voters must forecast the qualified region

The final row is the substantive trade-off, and it is a cost rather than a refinement. A runoff voter knows exactly which comparison she is participating in. An FVS voter allocates a scarce rejection budget before Q is determined, and may spend a rejection on a candidate who does not qualify or withhold one from an unexpected qualifier. **Section 8.2 names this the viability-visibility condition; it is the price FVS pays for collapsing two events into one.** The cost is heaviest under the avoidance reading of Section 5.5, which asks voters to condition explicitly on expected viability.

Section 3.3 records a second difference. A rank cutoff always produces a qualified set of the stipulated size; a quota produces however many clear the line.

4.3 The constructive vote of no confidence

The principle of Section 1.3 has an established constitutional analogue. Under Article 67 of the German Basic Law, and in systems modelled on it, a legislature may remove a head of government only by simultaneously electing a successor by majority. Opposition alone is insufficient; a replacement must be constituted in the same act.

The analogy clarifies what FVS does and does not require.

The parliamentary rule requires a replacement majority. The removing coalition must actually assemble a majority behind a named successor.

FVS requires only replacement qualification. An opponent of the sole qualifier needs a rival to clear the affirmative gate — that is, to attract more than α of issued authorizations. That is a much weaker demand than assembling a majority for it. **Once two candidates qualify, the contest between them is settled entirely by the negative channel: the winner is whichever carries the lower R_j, a quantity produced by rejections and wholly independent of how much affirmative support the challenger has beyond the threshold.** A challenger qualifying at 21% defeats a sole-qualifier incumbent at 60% whenever the challenger's recorded rejection is lower. The opposition may therefore prevail without ever having formed an affirmative majority for anything, and without the rejection ordering being constrained by relative affirmative strength.

FVS accordingly sits between unrestricted negative voting and the constructive vote of no confidence: more demanding than the former, considerably less demanding than the latter.

The comparison has an unflattering implication that should be stated. The parliamentary device operates where a successor is always available from a sitting chamber. **A presidential election has no sitting chamber from which to constitute a successor**, which is why candidate-set rejection in 2002 would have left France without a head of state, and why Section 9.5 cannot treat the continuation rule as an implementation detail.

4.4 Negative ballots and threshold aggregation

Negative voting is not new. Brams (1977) studies incentives to cast a negative vote when negative and positive votes enter a common total. Felsenthal (1989) combines approval with disapproval; Alcantud and Laruelle (2014) characterize dis&approval voting. Myerson (1999) compares positive and negative scoring variants. Kang (2010) develops a normative account of voting as veto. García-Lapresta, Marley, and Martínez-Panero (2010) characterize best-worst scoring systems. González, Laruelle, and Solal (2019) analyze approve/neutral/disapprove ballots. Gibbs and Chappell (2021) investigate multiple positive and negative votes; Jarabinský, Oreský, and Čech (2025) compare D21 variants. Foley (2024) discusses reforms involving affirmative screening.

The closest formal relative of the FVS selection stage is the threshold rule of Aleskerov, Yakuba, and Yuzbashev (2007), extended in Aleskerov, Chistyakov, and Kalyagin (2010). That literature aggregates three-graded evaluations and proves the threshold rule uniquely satisfies its axiom system. The rule is explicitly non-compensatory — a single lowest-grade evaluation is not offset by multiple highest-grade evaluations — and orders alternatives first by the count of lowest-grade evaluations.

Taken alone and applied to an already-fixed comparison set, the FVS selection stage is an instance of ordering by count of the negative grade, and that logic is characterized in the prior literature. The paper does not claim

it as new.

Three differences remain, and the case rests on them jointly.

1. **Bounded per-voter budgets.** In the threshold-aggregation setting a voter may assign the lowest grade to any number of alternatives. Section 2.2.2 argues that scarcity is what gives the negative channel its prioritization content, and also what creates the vulnerability of Section 9.3.2.
2. **An affirmative qualification gate preceding the negative ordering.** The threshold rule orders the whole supplied set; FVS first constructs Q. Theorem 1 concerns exactly this step.
3. **A bounded comparison set.** The gate and the affirmative budget jointly cap qualification at four and produce sole-qualifier decisiveness.

The paper's position is that (2) is substantive and that (1) and (3) follow from parameterization.

4.5 The architectural comparison

Rule family	Voter message	Winner logic	Distinction from FVS
Plurality	One positive mark	Most positive marks	No direct candidate-specific rejection
Two-round runoff	One mark per round	Rank cutoff, then majority	Two events; Q known at selection
Approval voting	Any approved subset	Most approvals	Positive bandwidth not scarce; no separate rejection
Combined approval–disapproval (compensatory)	Positive, neutral or negative per candidate	Common signed total	Signed information recombined on one scale
Threshold aggregation (Aleskerov et al.)	Three-graded evaluation of every alternative	Order by count of lowest grade	Unbounded grading; no affirmative gate
Lexicographic dis&approval (González et al.)	Approve / neutral / disapprove	Characterized lexicographic orders	No scarcity constraint; no affirmative gate
Best-worst / D21-type rules	Bounded positive and negative marks	Usually a common score	Signed information commonly recombined
One YES-or-NO negative vote	One positive or one negative mark	Highest net tally	Cannot transmit authorization and prioritized opposition together
FVS	At most one authorization and two rejections	Affirmative qualification, then minimum rejection	Scarce typed signals, non-compensation, bounded qualified set

Table 1. Nearby rules and the architectural distinction claimed for FVS.

4.6 Functional separation rather than motivational purity

FVS does not infer that a positive mark reflects enthusiasm or that a negative mark is free from instrumental calculation. A voter may authorize B chiefly to prevent A. The institution records the authorization because the voter chose an action that can help B acquire standing; it does not pretend to recover the motive.

5. Formal Definition of FVS

5.1 Domain and ballots

Let $C = \{1, \ldots, m\}$ be a frozen set of formally nominated candidates competing for one office, with $m \geq 4$. The family is FVS(α, d), with $\alpha \in (0, 1)$ a strict affirmative qualification threshold and $d \in \{1, \ldots, m - 1\}$ the maximum number of active rejections per voter. The specification studied is FVS(0.20, 2).

Each voter i submits $b_i = (p_i, N_i)$, where $p_i \in C \cup \{\varnothing\}$ is one authorized candidate or none, and $N_i \subseteq C$ satisfies $|N_i| \le d$ and $p_i \notin N_i$. Every mark is optional; a fully blank ballot is valid.

The candidate set is frozen when voting begins. Death, withdrawal or disqualification afterwards must be handled by a pre-specified rule rather than by silently deleting the candidate and recomputing denominators.

5.2 Counts and denominators

Let E be valid submitted ballots, T the number of voters issuing at least one candidate-directed signal, and V the number issuing a positive authorization, so $0 \le V \le T \le E$. For each candidate j:

$$A_j = |\{ i : p_i = j \}|, \; D_j = |\{ i : j \in N_i \}|.$$

Whenever $V > 0$, $S_j = A_j / V$. Whenever $T > 0$, $R_j = D_j / T$.

S_j measures concentration inside the affirmative pool; it is not a share of all participants. R_j measures the incidence of explicit rejection among active participants; it does not imply that $1 - R_j$ affirmatively accepts, knows, or has evaluated the candidate.

5.3 Qualification and selection

If $V > 0$, $Q = \{ j \in C : S_j > \alpha \}$. If $Q \neq \varnothing$, the winner is chosen from $W_1 = \arg\min_{j \in Q} R_j$, with ties resolved by $W_2 = \arg\max_{j \in W_1} S_j$, and any remaining tie by a predetermined candidate-symmetric neutral procedure specified in advance. Because every R_j shares denominator T, minimizing R_j is equivalent to minimizing D_j; the rate is retained for public interpretation across elections.

If $V = 0$ or $Q = \varnothing$, FVS returns **candidate-set rejection**. Section 9.5 treats the continuation rule.

5.4 The complete public output

Let F_B = (E, T, V, A, D, S, R, Q, w), with w replaced by the candidate-set-rejection symbol when no winner is produced. An implementation must publish, for every candidate, A_j, A_j / T, S_j, D_j and R_j, together with T/E, V/E and V/T. If a denominator is zero, the associated rate is reported as not applicable rather than imputed as zero.

This separates decision value from record value. A rejection of a candidate finishing just below the threshold does not affect the current winner ordering but remains part of the official negative record.

5.5 The required definition of the negative mark

The rule supports only these readings: **+j** means an affirmative authorization was issued; **−j** means an active rejection was issued within a bounded budget; **blank for j** means neither signal was allocated. It does not support the claims that +j identifies a private first choice, that −j identifies absolute unacceptability, or that blank means indifference or consent.

An adopting polity must additionally fix, in the ballot instructions, which of two readings the negative mark carries.

- **Aversion.** *Mark the candidates you most oppose.*
- **Avoidance of viable rivals.** *Mark the candidates you most want to prevent from winning.*

Both are legitimate, and the architecture of Section 2.2 is compatible with either. **They are not equivalent in effect, and the difference is not confined to wording.** The two readings generate a dilemma, stated here and demonstrated in Sections 6.7, 7.2 and Appendix C.

	Aversion	**Avoidance of viable rivals**
Absorbing entry (Appendix C)	**Exploitable**: an entrant with 2% of authorizations, incapable of qualifying, changes the winner	**Inert**: an entrant who cannot win attracts no rejection
Negative channel at $q \le d + 1$	Does independent work; R_j and S_j orderings differ	**Collapses**: $D_j = T - A_j$, so the rule selects the plurality winner within Q (Proposition 8)

	Aversion	Avoidance of viable rivals
Cognitive demand	Evaluate candidates	Evaluate candidates **and** forecast viability
Realized q in five French first rounds	—	q ≤ 3 in all five (Section 3.2), so the collapse applies to every observed case

Table 2. The two admissible readings of the negative mark.

Neither reading dominates. Aversion preserves the negative channel's independent content at the price of a cheap manipulation; avoidance closes the manipulation at the price of reducing the rule to plurality within Q in exactly the configurations that recent large-field elections produce. **The choice is constitutional, not editorial, and this paper does not resolve it.** Section 10.4 proposes it as a primary experimental question; Section 10.6 proposes the entry game that would establish how robustly avoidance closes the manipulation outside the configuration of Appendix C.

5.6 No validity threshold on V/T

This paper specifies no acceptable range for V/T, and the omission is a position rather than a gap.

Any such range would make the result conditional on how many participants declined to issue an authorization — that is, it would read positive abstention as collective negation. Section 2.6 declines to make that inference about an individual ballot. **The inference does not become better grounded by being made in aggregate.**

Two consequences follow.

The mechanism's defensibility does not vary with V/T. At any V/T the rule does what it claims: record which proposals were authorized, which were rejected, and by how many. The 99-to-1 case is not a malfunction; it is an accurate report that one participant supplied a direction and ninety-nine did not.

The winner's mandate strength does vary with V/T, and that is what publication is for. Because the design declines to protect interpretation through a validity condition, it must protect it through a reporting condition: A_j**/T and V/T are published inseparably from the result, with no triggering level, and any mandate claim must be stated relative to them.**

6. Structural Properties

6.1 Affirmative normalization and the proposal cap

Lemma 1. If V > 0, then $\sum_{j \in C} S_j = 1$.

Proof. Every member of the affirmative pool issues exactly one authorization, so $\sum_j A_j = V$. ∎

Proposition 1. Under FVS(α, d), q = |Q| satisfies q < 1/α. If $\alpha \geq 1/(d+3)$, then q ≤ d + 2. Under FVS(0.20, 2), at most four candidates qualify.

Proof. If q candidates qualify their shares sum to more than qα; Lemma 1 gives qα < 1. If $\alpha \geq 1/(d+3)$ then q < d + 3. ∎

The strict inequality matters: five candidates can each receive exactly 20%, but none qualifies under $S_j > 0.20$.

The derivation is an upper-bound derivation. The interval 0.20 ≤ α < 0.25 is precisely the set of thresholds guaranteeing a cap of four while permitting four to qualify, and 0.20 is its lowest member. It does not independently establish that a candidate at 19.58% is politically distinguishable from one at 20.01%. The threshold's role as an exclusion floor is a consequence of an upper-bound derivation, and should be understood as a public line fixed in advance rather than a substantive measurement. Section 9.5 draws a constraint from this for continuation rules.

6.2 Threshold–budget compatibility

Corollary 2. If $\alpha \geq 1/(d+3)$ and a voter authorizes one member of Q, that voter can reject all other qualified candidates except at most one.

Proof. Proposition 1 gives q ≤ d + 2; a voter authorizing one qualified candidate faces at most q − 1 ≤ d + 1 rivals and can reject d. ∎

The corollary is tight only at the cap: when q < d + 2 the budget covers every qualified rival. Under FVS(0.20, 2) the guarantee of one unmarked qualified rival holds at q = 4 and not at q ≤ 3. Proposition 8 draws the consequence.

6.3 Directional abstention neutrality

Proposition 2. Hold every A_j fixed. Adding voters who issue no positive authorization leaves V and every S_j unchanged. Such voters can still alter D, T and the rejection ordering.

Proof. A voter enters V only by issuing a positive mark. ∎

Random reduction of V barely changes Q, because random subsampling of the affirmative pool preserves shares approximately. The sole-qualifier case therefore requires *structured* affirmative abstention or structured fragmentation, not merely low affirmative turnout. Section 9.3 distinguishes the two.

6.4 Stagewise non-compensation

Proposition 3. Rejections do not affect membership in Q. For qualified a and b, if $R_a < R_b$ then b cannot defeat a merely by having higher S_b.

Corollary 3 (sole-qualifier decisiveness). If Q = {a}, then w = a regardless of R_a.

Negative signals can rank multiple affirmatively standing proposals; they cannot create a rival that no sufficient positive coalition authorized.

6.5 The aggregate rejection bound

Proposition 4. For FVS(α, d) with T > 0, $\sum_{j \in C} R_j \leq d$, with equality if and only if every signal-sending participant uses the full rejection budget.

Proof. Each active voter issues at most d rejections, so $\sum_j D_j \leq dT$. ∎

This is a capacity bound, not conservation in general. Lemma 1 and Proposition 4 jointly define bounded expression here: the affirmative channel is normalized to one, the negative channel capped at d.

6.6 Own-signal monotonicity

Proposition 5. Suppose a wins. A legal ballot change adding an authorization for a, or transferring one to a, holding all other candidate-directed marks fixed, cannot make a lose.

Proof. If V is unchanged, S_a rises and a rival's share falls. If a previously non-authorizing ballot adds +a, then $(A_a + 1)/(V + 1) \geq A_a / V$ since $V \geq A_a$, while rivals' shares weakly fall. Any increase in T rescales all rejection rates identically. ∎

Proposition 6. Suppose a wins. Removing one rejection of a, holding every other candidate-directed mark fixed, cannot make a lose.

Proof. If the voter remains active, D_a falls and nothing else changes. If removal blanks the ballot, D_a and T each fall by one while rival counts are unchanged. ∎

These are mechanical, ceteris paribus properties, not equilibrium claims.

6.7 Bounded rejection, and what each reading of the negative mark implies

The argument for boundedness is in Section 2.2.2 and depends on no behavioural assumption: prioritization carries information only when the budget is scarce. Two propositions describe what the bounded channel then does, and they correspond to the two readings of Section 5.5.

Proposition 7 (complement collapse under an exhaustive profile). Suppose every active voter issues one authorization and rejects every non-authorized candidate. Then $D_j = \mathrm{T} - A_j$ for every j, and within any fixed Q, minimizing D_j is equivalent to maximizing A_j.

Proof. Every voter either authorizes j or rejects j. ∎

Proposition 7 concerns a strategy profile that no reading of the mark prescribes, and the paper does not claim it is dominant. The next result is stronger, because the avoidance reading prescribes the relevant behaviour directly.

Proposition 8 (collapse under the avoidance reading). Under the avoidance reading of Section 5.5, suppose every participant authorizes a qualified candidate and $\mathrm{q} \le \mathrm{d} + 1$. Then each voter rejects every qualified rival, so $D_j = \mathrm{T} - A_j$ for every $\mathrm{j} \in \mathrm{Q}$, and the winner is the plurality leader within Q.

Proof. A voter authorizing $\mathrm{a} \in \mathrm{Q}$ faces $\mathrm{q} - 1 \le \mathrm{d}$ qualified rivals. The avoidance reading directs the rejection budget at the qualified rivals the voter most wants to prevent from winning; since the budget covers all of them, all are rejected. Hence D_j counts every participant not authorizing j, which is $\mathrm{T} - A_j$. Minimizing D_j over Q is maximizing A_j over Q. ∎

If some participants authorize outside Q, each such voter faces $\mathrm{q} \le \mathrm{d} + 2$ rivals and must spare at least one, so $D_j = (\mathrm{T} - A_j) - s_j$ where s_j counts the non-Q authorizers sparing j. The rule is then plurality within Q corrected by a second-preference term from those voters alone.

Under FVS(0.20, 2), Proposition 8 applies whenever $\mathrm{q} \le 3$, and Section 3.2 reports $\mathrm{q} \le 3$ in all five French first rounds examined. The consequence must be stated plainly: **if a polity adopts the avoidance reading, the negative channel does no independent selection work in the configurations that recent large-field elections produce.** The gate still does work — it determines Q, and Corollary 3 still binds at $\mathrm{q} = 1$ — but the selection stage reduces to plurality within the qualified set.

The rationale for d = 2 must be qualified accordingly. For a maximum qualified set of four, one rejection may be too narrow to identify more than one dangerous alternative, while three would allow a voter authorizing A to reject B, C and D. Conditional on permitting more than one defensive target while requiring at least one other qualified candidate to remain unmarked at the cap, two is the unique integer budget. But that objective is stipulated rather than derived; the property holds only at $\mathrm{q} = 4$; and $\mathrm{q} = 4$ did not occur in any of the five elections of Section 3.2. **The honest statement is that d = 2 is a defensible companion budget for a four-candidate field with a filled qualified set, not a parameter derived in general.**

6.8 Ballot-space complexity

With m candidates and d = 2, the number of legal ballots is

$$\mathrm{B}(\mathrm{m}) = \sum_{k=0}^{2} \binom{m}{k} + m \sum_{k=0}^{2} \binom{m-1}{k} = (m^3 + 3m + 2)/2.$$

For m = 4, B(4) = 39, about $\log_2 39 \approx 5.29$ bits. This shows only that the ballot distinguishes a substantial set of typed public messages without requiring a complete ranking.

6.9 Entry, positive-pool externality, and reinforcement

> **Proposition 9 (zero-uptake entry neutrality, ceteris paribus).** Add a candidate x with $A_x = 0$, holding all marks directed at existing candidates fixed. The winner among existing candidates is unchanged.

The ceteris paribus clause holds fixed the marks directed at existing candidates, which is precisely the channel through which entry operates. Appendix C gives a closed-form case in which an entrant holding 2% of authorizations, incapable of qualifying, changes the winner under the aversion reading.

Entry with positive uptake is not neutral either. Because $S_j = A_j$ / V, a new positive mark enlarges the pool and reduces every other candidate's share. For $\alpha = 0.20$, define $L_j = 5\ A_j - \mathrm{V}$; candidate j qualifies exactly when $L_j > 0$. An additional authorization for j changes L_j by +4; one for another candidate by −1; a transfer to j by +5.

The sharpest illustration uses five ballots. Four voters authorize A, B, C and D; each has 25% and all qualify. A fifth authorizes E; all five then have exactly 20%, so Q = ∅.

FVS is homogeneous, anonymous, and candidate-neutral if the final tie-break is neutral. It is not globally reinforcing: a candidate can win in two electorates and lose in their union, because the qualified set changes.

7. Strategic and Axiomatic Stress Tests

7.1 No privileged definition of sincerity

FVS specifies legal public acts but not a unique mapping from private preferences to them. Testing majority, Condorcet or participation criteria requires an explicit mapping, and results under one mapping are not unconditional properties of the rule. General results on manipulation provide no basis for claiming strategy-proofness (Gibbard, 1973), and FVS is manipulable under both readings.

This paper does not claim that expending the full rejection budget is weakly dominant, and it is not. A voter who dislikes C somewhat and D greatly can be made worse off by adding −C, if that mark moves C above D in the final ordering.

7.2 Two analytical mappings, and why the results diverge

The stress tests below use two mappings, one for each reading of Section 5.5. For four candidates and a strict ranking $x_1 \succ x_2 \succ x_3 \succ x_4$:

Aversion mapping: $(+x_1, -x_3, -x_4)$ — authorize the top-ranked candidate, reject the two lowest-ranked candidates in the field. Under this mapping, with V = T = N, letting F_j be j's first-place share and U_j the share ranking j in the top two:

> $S_j = F_j$, $R_j = 1 - U_j$.

FVS(0.20, 2) under the aversion mapping therefore qualifies candidates with more than 20% first-place support and selects, among them, the candidate with greatest top-two coverage.

Avoidance mapping: authorize the top-ranked candidate; reject the d lowest-ranked members of Q {authorized candidate}. By Proposition 8, at q ≤ 3 with universal in-Q authorization this reduces to plurality within Q.

Because the two mappings differ, several standard counterexamples are reading-specific. This must be reported rather than suppressed: a paper that presents one mapping's failures as properties of the rule misstates what an adopting polity would obtain.

Criterion	Aversion mapping	Avoidance mapping	Depends on the reading?
Pareto efficiency	Satisfied	Satisfied	No
Majority criterion	**Fails** (§7.3)	Satisfied in the §7.3 profile; by Prop. 8 a qualified plurality majority wins at q ≤ 3	**Yes**

Criterion	Aversion mapping	Avoidance mapping	Depends on the reading?
Condorcet consistency	Fails	Fails	No
Condorcet-loser exclusion	**Fails** (§7.4)	**Fails** — same profile, Q is a singleton and no rejection is decisive	No
Participation / no-show	**Fails** (§7.5)	Does not fail in the §7.5 profile; failures arise through Q-dependence, which is reading-independent	**Partly**
Strategy-proofness	**Fails** (§7.6)	Fails; the §7.6 deviation is not profitable, but plurality-style manipulation applies at q ≤ 3	Profile-specific

Table 3. Stress tests under both mappings. Entries marked "satisfied in the profile" record only that the stated counterexample does not transfer, not that the criterion holds in general.

Two observations follow. **The Condorcet-loser failure is the most robust**, because it turns on a singleton qualified set in which no rejection can be decisive; it survives both readings and is therefore the sharpest cost of Corollary 3. **The majority failure is the least robust**, because it depends on the aversion reading's top-two coverage statistic; under avoidance at q ≤ 3 the rule inherits plurality's satisfaction of the majority criterion within Q.

7.3 Majority failure (aversion mapping)

Consider 51 voters with A ≻ B ≻ C ≻ D and 49 with B ≻ C ≻ A ≻ D. Both A and B qualify ($S_A = 0.51$, $S_B = 0.49$).

Under the aversion mapping the 51 reject {C, D} and the 49 reject {A, D}, so $D_A = 49$ and $D_B = 0$. Candidate B wins although A holds a first-preference majority and is the Condorcet winner.

Under the avoidance mapping each voter rejects their single qualified rival: $D_A = 49$, $D_B = 51$, and A wins. The failure does not transfer.

7.4 Condorcet-loser failure (both mappings)

Take five voters with B ≻ A ≻ D ≻ C; B ≻ A ≻ C ≻ D; C ≻ D ≻ A ≻ B; D ≻ C ≻ A ≻ B; A ≻ C ≻ D ≻ B. First-place shares are B = 40% and A = C = D = 20%, so under the strict threshold Q = {B}.

Candidate B wins by Corollary 3 despite losing pairwise to A, C and D. **No rejection is decisive in a singleton qualified set, so the failure is identical under both readings.** This is the sharpest cost of constructive responsibility and the paper does not soften it.

7.5 Participation failure (aversion mapping)

With four voters ranked A ≻ B ≻ D ≻ C; D ≻ A ≻ B ≻ C; D ≻ A ≻ B ≻ C; B ≻ C ≻ D ≻ A, the aversion mapping gives A = 25%, B = 25%, D = 50%, all qualifying; A receives one rejection, B and D two each, so A wins. Adding a fifth voter with C ≻ B ≻ A ≻ D casting (+C, −A, −D) makes shares A = B = C = 20% and D = 40%, so Q = {D}. The entering voter ranks A above D and is worse off by participating.

Under the avoidance mapping the original four-voter profile already elects D, so this particular deviation produces no harm. Participation failures nonetheless remain available under either reading, because they operate through the dependence of Q on V, which is reading-independent.

7.6 A profitable positive reallocation (aversion mapping)

Three voters ranked C ≻ D ≻ B ≻ A; A ≻ D ≻ B ≻ C; A ≻ D ≻ B ≻ C give Q = {A, C}; under aversion A has one rejection and C two, so A wins. The first voter can retain (−B, −A) while moving the positive mark from C to D. Then Q = {A, D}, D has no rejection and wins; since the first voter prefers D to A, the deviation is profitable.

Under the avoidance mapping the deviation is unprofitable in this profile, because each voter rejects only qualified rivals and the counts are unchanged. Manipulability is not thereby removed: at $q \leq 3$ the avoidance reading inherits plurality's standard manipulations, in which a voter whose preferred candidate cannot lead authorizes a different qualified candidate.

7.7 Reinforcement and field dependence

In electorate 1, ballots −A and +A give Q = {A} and A wins. In electorate 2, ballots +A, +A, +B qualify both; rejection is tied and A wins the affirmative tie-break. In the union, A has three authorizations and B one, both qualify, but A carries one rejection and B none; B wins. The failure operates through the qualified set and is reading-independent.

8. Information, Viability, and the Public Record

8.1 Qualification is not an awareness test

Suppose 100 voters each issue one positive mark. Twenty-one know and authorize A; 79 authorize B and have never encountered A. If the 21 reject B and the rest do not reject A because they do not know A, then $S_A = 0.21$, $S_B = 0.79$, $R_A = 0$, $R_B = 0.21$. Both qualify and A wins.

The gate has prevented a candidate with no positive base from exploiting obscurity, but has not guaranteed that low rejection reflects informed tolerance. It raises the cost of exploiting unfamiliarity; it does not close the channel. The divergence is sharper when V/T is low, and neither S_j nor A_j / T measures awareness among non-supporters; Section 10.3 proposes an independent recognition measure.

8.2 The viability-visibility condition

Campaigns, debates, endorsements and polls typically give voters an approximate sense of which candidates may approach the threshold. This practical focusing is not generated by the tally formula; polls can be wrong, manipulated, or self-fulfilling.

The mechanism relies on a **viability-visibility condition**: candidates plausibly near qualification should be sufficiently visible for active rejection to be meaningful. When it fails, negative marks may be allocated to non-qualifiers or withheld from an unexpected qualifier, and low rejection can partly reward obscurity.

This condition is the price of collapsing a two-event sequence into one ballot, and **it bears unequally on the two readings of Section 5.5.** The aversion reading asks voters only to evaluate candidates; the avoidance reading asks them to evaluate candidates and forecast viability, and therefore depends on the condition directly. A polity choosing avoidance to close the manipulation of Appendix C thereby increases the rule's dependence on the quality of public viability signals and on the actors who supply them.

8.3 Near-threshold speech is not erased

A candidate at 19.9% and one at 20.1% can be politically similar even though only one has standing. The threshold must be discontinuous for decision; the public record need not be discontinuous in meaning. Both candidates' totals remain in F_B and may inform future nominations and subsequent elections.

8.4 What the winner's mandate means

An FVS result supports a compound and limited statement:

> The winner acquired sufficient concentration among issued affirmative authorizations to enter the qualified set and then received the least recorded active rejection among candidates in that set.

It does not support the claims that the winner had majority approval, that unmarked voters found the winner acceptable, that low rejection proves familiarity, or that the winner was the Condorcet choice. Any mandate claim must be stated relative to the published V/T and A_j/T.

The 99-to-1 sole-qualifier case is procedurally authorized but substantively weak, and its record would show one authorization against 99 rejections. The contrast with a second-round figure of 82.21%, which records no such thing, is the point of Section 3.4.

8.5 Polarity fidelity, not complete semantic fidelity

FVS preserves the sign and institutional function of marks voters actually issue. It cannot preserve reasons, intensity, conditionality or unexpressed rankings. The defensible claim is **bounded polarity fidelity**: under a severe message budget, a negative candidate-directed act remains negative, and positive and negative totals are not collapsed algebraically.

9. Normative Trade-offs and Institutional Limits

9.1 Constructive minorities and negative majorities

FVS gives organized affirmative minorities agenda power. A small group supplying the only qualified proposal can prevail over a much larger group that agrees only in opposition. The defence rests on Section 2.3; Section 4.3 gives the institutional analogue and notes that FVS demands strictly less than the parliamentary version.

The cost is that a highly rejected sole qualifier can take office, and Section 7.4 shows this includes a Condorcet loser under both readings of the negative mark. If a polity believes explicit majority rejection must always prevent election, FVS is the wrong rule unless supplemented by a rejection ceiling or a status quo that can itself qualify. Such an amendment changes the principle and should be analyzed as a different mechanism.

9.2 Positive abstention is free but not costless

No voter is compelled to issue a positive mark, but declining to enter V relinquishes direct influence over the distribution of affirmative standing. This is the endogenous consequence of declining to perform the act that creates standing.

9.3 Fragmentation and two directions of proliferation

The threshold operates on candidates, not on ideological families inferred by the counting authority. A tendency commanding 54% of issued authorizations divided equally among three candidates at 18% qualifies nobody while a rival at 46% qualifies.

Two structurally different profiles must be distinguished. In the **99-to-1 profile**, the large group issues no authorization at all, and "supplied no direction" is literally accurate. In the **18-18-18 profile**, the large group issues authorizations and directs them at candidates of one tendency; describing them as having supplied no direction would be inaccurate. What the rule asserts is narrower: **they supplied direction but did not concentrate it into a qualifying candidacy, and the rule requires formal concentration.**

The demand is legitimate and the rule cannot relieve it without transferring proposal power to officials or becoming a different mechanism. It is nonetheless a real burden falling hardest on the most pluralistic tendencies. A comparable outcome obtains under plurality.

9.3.1 Affirmative proliferation is costly

A faction with share f cannot, without outside support, qualify k candidates unless $f > k\alpha$. Under $\alpha = 0.20$ this requires more than 40%, 60% or 80% for two, three or four.

9.3.2 Negative proliferation may be free, and whether it is depends on Section 5.5

A candidate entered solely to absorb rejection need not qualify. Because the negative budget is bounded, each voter's rejections go to a limited priority set, and an entrant who enters that set displaces someone from it.

Appendix C gives a closed-form instance. Three qualified candidates hold 40%, 24% and 36% of authorizations. Without an entrant, the leading candidate wins with a rejection rate of 3/5. An entrant positioned outside the field, attracting 2% of authorizations and therefore incapable of qualifying, becomes the most-opposed candidate for two of the three voter blocs under the aversion reading. Those blocs redirect one rejection each to the entrant, sparing the centre candidate, whose rejection rate falls from 19/25 to 1/50. **The centre candidate, holding the smallest affirmative share of the three, wins.** Under the avoidance reading the entrant cannot win, attracts no rejection, and every rejection total is unchanged; the original winner prevails.

The exploitability of negative proliferation is therefore not a fixed property of FVS. It is a property of the definition the adopting polity writes into its ballot instructions.

This is one reason to prefer the avoidance reading. Section 5.5 records the reason against it: by Proposition 8, avoidance reduces the rule to plurality within Q whenever $q \leq 3$. **A polity cannot obtain both an independent negative channel and immunity to absorbing entry from a bounded rejection budget alone.** The closed-form example establishes that the manipulation exists under aversion in at least one configuration; Section 10.6 states the entry game required to determine how general it is.

9.4 Negative marks are comparative, not vetoes

With at most two rejections and as many as four qualified candidates, an unmarked qualified candidate may be disliked, unknown, neutral or simply lower priority. The rule must say "recorded rejection" rather than "unacceptability" unless a separate behavioural model justifies the stronger term.

9.5 Candidate-set rejection: three constraints and an analytical baseline

If $V = 0$ or $Q = \varnothing$, the rule returns the candidate set rather than a winner. **By Section 2.3 this is a choice of reversion outcome, and reversion position determines proposal power.** By Section 3.2 the choice is not hypothetical.

This paper does not select a continuation rule as a policy recommendation. That choice concerns eligibility for public office and the timing of national elections; it belongs to a constitution rather than to a counting rule. What the paper supplies are constraints that follow from its own arguments, and — because a running case cannot be left unresolved — a baseline used for analysis.

Constraint 1: the reversion must not be returned to the agenda setter. Automatic incumbent retention or automatic status quo assigns the reversion to the actor whose advantage Section 2.3 was invoked to explain. Applied to 2002 it would have continued in office a President who had just received 19.88%. Excluded.

Constraint 2: an exclusion threshold, if adopted, must not equal α. A continuation rule may bar some candidates from the subsequent contest, and doing so is the only element of the design capable of imposing a cost on the proposing side. But Section 6.1 concedes that α is derived as an upper bound and that its position is a convention. **A line the paper declines to treat as substantive cannot be used to terminate a political career.** Applied to 2002, excluding every candidate who failed α would have barred a sitting President and a sitting Prime Minister from national candidacy on a shortfall of 0.12 percentage points.

Constraint 3: a termination condition is required. Repeated candidate-set rejection under a shrinking eligible pool must not recurse indefinitely.

Analytical baseline. Where this paper needs a determinate continuation — principally for the 2002 case and for the modelling proposed in Section 10.7 — it uses: *a re-contest after a fixed interval, with candidates whose affirmative share fell below a low non-starter floor $\beta \ll \alpha$ barred from the re-contest, and with a pre-specified fallback rule invoked if the re-contest also returns candidate-set rejection.* The baseline satisfies all three constraints. It is a modelling choice, not a recommendation, and alternatives within the admissible class belong in robustness analysis.

Applied to 2002 with $\beta = 0.05$, the baseline bars the nine candidates below 5% who together absorbed roughly a quarter of first preferences, and re-contests among the remaining seven. It addresses the proximate cause of the dispersion identified in Section 3.1 without barring any candidate who approached viability. Its cost is that serious candidates face a weaker incentive to avoid splitting, since they expect to survive; Section 10.7 proposes measuring that effect against a higher floor.

9.6 Scope of the parameterization

The parameterization applies only to single-winner elections with at least four candidates and to electorates in which affirmative authorization can concentrate above the threshold. Multi-member elections, party lists, proportional representation, write-ins, recall elections and referenda require separate specifications.

Voters must understand that the positive denominator excludes positive abstainers, that negative marks are bounded and comparative, that the negative mark carries whichever reading of Section 5.5 the polity adopted, and that a sole qualifier wins regardless of rejection. Ballot instructions and result reporting are part of the mechanism.

9.7 No mechanical claim about negative campaigning

Affirmative qualification does not prove that negative campaigning cannot improve an attacker's positive share. If A has 30 authorizations and B has 70, S_A = 30%; if an attack causes 20 B supporters to issue no positive mark while A's count is unchanged, S_A = 30/80 = 37.5%. Only under a restrictive campaign technology is qualification mechanically closed to attack.

9.8 Upstream inequality remains

FVS regulates $C \to Q \to w$, not $\Omega \to C$. Parties, finance, media, ballot-access law, recruitment and incumbency still shape which proposals become candidates. Section 2.8 identified the distributive consequence of the affirmative default: the architecture disadvantages new entrants and proposals whose support is real but not yet widely encountered, and that disadvantage compounds with upstream inequality rather than offsetting it.

The one exception is the continuation rule of Section 9.5, which is the only element of the design capable of imposing a cost on the proposing side.

10. Research Agenda

10.1 The frequency of qualified-set sizes

Section 3.2 reports |Q| for five French presidential first rounds. Five observations from one country, one nomination system and one first-round rule support an illustration, not the scope claim of Section 1.4. Apply the same mapping to first-round or single-round multi-candidate presidential elections across a substantially larger set of countries and cycles, and report the frequency distribution of q ∈ {0, 1, 2, 3, 4}.

Two of the paper's arguments depend on the result. **The scope condition** — that concentration rather than field size is binding — becomes evidence rather than assertion only with this table. And **the frequency of q = 4** determines the standing of the d = 2 rationale of Section 6.7 and the reach of Proposition 8. If q = 4 proves rare across systems, the structural argument for d = 2 should be withdrawn rather than qualified, and the avoidance reading's collapse should be described as the normal case rather than a common one.

10.2 The semantics of the negative mark

Section 5.5 requires the polity to fix whether a rejection records aversion or avoidance of viable rivals, and Sections 6.7 and 9.3.2 show that each reading carries a distinct and serious cost. Run both instruction sets as experimental treatments with human subjects. Measure comprehension, allocation patterns, whether voters under the avoidance instruction in fact condition on viability, winner divergence, and susceptibility to absorbing entrants. **This is the paper's primary open empirical question**, and the dilemma of Section 5.5 cannot be adjudicated without it.

10.3 Recognition, independently measured

Neither S_j nor A_j / T measures awareness among non-supporters. Construct an independent recognition measure and estimate the joint distribution of recognition and R_j. This is required before any claim about informed tolerance can be assessed.

10.4 V/T as an open measurement

Section 2.2.3 reports evidence that voters answer the supply question at scale when an event is framed as a qualification contest. It does not establish what happens when declining to authorize is a distinct available option. Measure the distribution of V/T across settings, instructions, field sizes and stakes. By Section 5.6 this is a measurement, not a validity test.

10.5 The qualification gate and parameter sensitivity

Compare a ternary ballot aggregated without a gate against FVS, holding ballot and candidate information constant. Measure how often the gate changes the winner, how often q = 1, and how often Q = ∅. Vary α and d jointly across $\alpha \geq 1/(d+3)$. Study behaviour near the threshold: whether voters can identify the qualification region, and whether candidates just below it differ systematically from those just above.

10.6 The entry game

Appendix C establishes that absorbing entry changes the winner under the aversion reading in one configuration, and is inert under the avoidance reading in that same configuration. Neither result is general. Build a full entry game: candidates choose whether to enter, merge or withdraw; voters respond to expected viability; both readings are compared; positions and shares range over a specified space. Determine whether absorbing entry is profitable in equilibrium, for which factions, and whether the avoidance reading closes it robustly or only where the entrant is sufficiently peripheral.

10.7 Continuation rules

Model the election under continuation rules satisfying the three constraints of Section 9.5, varying the non-starter floor β and measuring the incentive to engineer $Q = \varnothing$, the effect on pre-election coordination, and termination behaviour under repeated failure. The 2002 configuration provides a calibration target, and the baseline of Section 9.5 provides a reference point.

10.8 Campaign effects

Specify whether campaign expenditure changes positive counts, negative counts, abstention, awareness, or several channels at once. Only then compare private and field-wide returns to negative campaigning across plurality, net signed rules and FVS.

10.9 Public-record value and pilots

Pilot in organizations, associations, student bodies and online communities. Publish anonymized F_B outputs, completion rates, V/T, qualification failures and recountable code. Test whether observers interpret mandates more accurately when positive and negative records are separated.

11. Conclusion

Electoral institutions pose two questions. *Which proposals should be under consideration?* asks voters to supply direction. *Among those under consideration, which should be avoided?* asks them to evaluate what has been supplied. Two-round systems pose the questions in sequence; the French first round of 2002 and the runoff three weeks later are a clean instance, with the same electorate answering both. FVS poses both on one ballot and separates them by channel rather than by event.

That choice determines the architecture. Selection is assigned to the negative channel because it answers the evaluation question directly. The channel is bounded because avoidance across several options is a prioritization problem, and prioritization carries information only when the budget is scarce. Qualification is assigned to the affirmative channel because the alternative is to let standing arise from the absence of objection; Theorem 1 shows that within a bounded three-state ballot imputing nothing to silence, that alternative is unavailable. The theorem is conditional on its design class and is not a general impossibility result.

The paper's sharpest finding is negative and concerns the same bounded budget that makes the negative channel informative. Because the budget is bounded, the positions in a voter's rejection priority set are a rival good, and an entrant who occupies one displaces someone from it. Whether this is exploitable depends on a definition the adopting polity must write into its ballot instructions. Under the aversion reading, a candidate holding two per cent of authorizations and incapable of qualifying can change the winner; Appendix C gives the arithmetic. Under the avoidance reading that manipulation is inert, but the rule reduces to plurality within the qualified set whenever four or fewer candidates qualify and voters authorize inside it — and in five French first rounds the qualified set never exceeded three. **A polity cannot obtain both an independent negative channel and immunity to absorbing entry from a bounded rejection budget alone.** This dilemma is stated rather than resolved, and Section 10.2 gives the experiment that would inform the choice.

The assignment of channels does not eliminate the difficulty that motivates it. An affirmative screen resolves unfamiliarity as non-qualification rather than as clearance — a choice between two errors, defended on grounds of remediability, legibility in the record, and what admission to the qualified set commits the electorate to. Its cost falls on new entrants and compounds with upstream agenda inequality.

The 2002 French election shows both halves of the case. Applying the gate to that year's first-preference distribution returns no qualifier, symmetrically, against an incumbent President, an incumbent Prime Minister and a far-right

challenger alike; the salient feature is that roughly a quarter of authorizations went to candidates below 5%, not that the leader fell 0.12 points short. The second round produced a figure of 82.21% that recorded almost none of what it appeared to record, from a comparison set the rank cutoff had generated rather than the electorate endorsed. FVS would not have prevented the dispersion and would not have produced a President. What it declines to do is package a dispersed affirmative distribution as an overwhelming mandate.

The architecture rests on a contestable premise. Opposition without a qualified alternative is politically meaningful but not sufficient to govern. A sole qualified candidate therefore wins even under overwhelming rejection, and Section 7.4 shows this includes a Condorcet loser under both readings of the negative mark. The burden this places on oppositions is real, falls hardest on pluralistic ones, and cannot be lightened without transferring proposal power to officials or adopting a different mechanism.

The 20% threshold should not be mythologized. It is derived as an upper-bound parameter capping the qualified set at four, and its role as an exclusion floor is a consequence of that derivation rather than an independent judgment. Correspondingly, no validity threshold is placed on V/T, because any such threshold would read positive abstention as collective negation.

FVS is therefore a bounded public-action language and a stagewise counting rule in which authorization, opposition and omission are typed differently; same-type signals aggregate together; unlike signals do not cancel; and negative judgment cannot create a positive proposal. Its selection logic has a characterized relative in the threshold-aggregation literature, and its claim rests on the gate that precedes it, on the theorem motivating the gate, and on the distinction between two institutional questions that motivates the theorem. Whether the resulting trade-off improves actual elections is a question for the cross-national tabulation of Section 10.1, the semantic experiment of Section 10.2, and the entry game of Section 10.6, rather than for assertion.

References

Alcantud, J. C. R., & Laruelle, A. (2014). Dis&approval voting: A characterization. *Social Choice and Welfare, 43*(1), 1–10.

Aleskerov, F., Chistyakov, V. V., & Kalyagin, V. (2010). Social threshold aggregations. *Social Choice and Welfare, 35*(4), 627–646.

Aleskerov, F., Yakuba, V., & Yuzbashev, D. (2007). A "threshold aggregation" of three-graded rankings. *Mathematical Social Sciences, 53*(1), 106–110.

Alós-Ferrer, C. (2006). A simple characterization of approval voting. *Social Choice and Welfare, 27*(3), 621–625.

Alós-Ferrer, C., & Granić, Đ.-G. (2012). Two field experiments on approval voting in Germany. *Social Choice and Welfare, 39*(1), 171–205.

Aragones, E., Gilboa, I., & Weiss, A. (2011). Making statements and approval voting. *Theory and Decision, 71*, 461–472.

Arrow, K. J. (1951). *Social Choice and Individual Values*. Yale University Press.

Balinski, M., & Laraki, R. (2011). *Majority Judgment: Measuring, Ranking, and Electing*. MIT Press.

Baujard, A., & Igersheim, H. (2011). Framed-field experiment on approval voting and evaluation voting. In B. Dolez, B. Grofman, & A. Laurent (Eds.), *In Situ and Laboratory Experiments on Electoral Law Reform* (pp. 69–89). Springer.

Baujard, A., Igersheim, H., Lebon, I., Gavrel, F., & Laslier, J.-F. (2014). Who's favored by evaluative voting? An experiment conducted during the 2012 French presidential election. *Electoral Studies, 34*, 131–145.

Brams, S. J. (1977). When is it advantageous to cast a negative vote? In R. Henn & O. Moeschlin (Eds.), *Mathematical Economics and Game Theory* (pp. 564–572). Springer.

Brams, S. J., & Fishburn, P. C. (1978). Approval voting. *American Political Science Review, 72*(3), 831–847.

Felsenthal, D. S. (1989). On combining approval with disapproval voting. *Behavioral Science, 34*(1), 53–60.

Foley, E. B. (2024). The nomination and election of statewide candidates. *University of Illinois Law Review, 2024*, 1607–1653.

García-Lapresta, J. L., Marley, A. A. J., & Martínez-Panero, M. (2010). Characterizing best-worst voting systems in the scoring context. *Social Choice and Welfare, 34*(3), 487–496.

Gibbard, A. (1973). Manipulation of voting schemes: A general result. *Econometrica, 41*(4), 587–601.

Gibbs, T., & Chappell, H. W., Jr. (2021). Elections with multiple positive and negative votes. *Homo Oeconomicus, 38*(1), 37–47.

González, S., Laruelle, A., & Solal, P. (2019). Dilemma with approval and disapproval votes. *Social Choice and Welfare, 53*(3), 497–517.

Jarabinský, I., Oreský, J., & Čech, P. (2025). D21 method: A path toward balanced elections? *Election Law Journal, 24*(4), 322–341.

Kang, M. S. (2010). Voting as veto. *Michigan Law Review, 108*(7), 1221–1282.

Kruger, J., & Sanver, M. R. (2021). An Arrovian impossibility in combining ranking and evaluation. *Social Choice and Welfare, 57*(3), 535–555.

Laruelle, A. (2021). "Not this one": Experimental use of the approval and disapproval ballot. *Homo Oeconomicus, 38*(1–4), 15–28.

Myerson, R. B. (1999). Theoretical comparisons of electoral systems. *European Economic Review, 43*(4–6), 671–697.

Romer, T., & Rosenthal, H. (1978). Political resource allocation, controlled agendas, and the status quo. *Public Choice, 33*(4), 27–43.

Tideman, T. N. (1987). Independence of clones as a criterion for voting rules. *Social Choice and Welfare, 4*(3), 185–206.

Appendix A. Compact Executable Specification

- **Domain.** One office; a frozen set C of at least four candidates; an electorate in which affirmative authorization can concentrate above the threshold.
- **Ballot.** Authorize zero or one candidate; reject zero, one or two different candidates; may not authorize and reject the same candidate.
- **Timing.** All marks submitted simultaneously; no intermediate qualification tally is revealed.
- **Counts.** E = valid ballots; T = ballots with at least one candidate-directed mark; V = ballots issuing an authorization. A blank ballot belongs to E but not T.
- **Shares.** If V > 0, $S_j = A_j$ / V. If T > 0, $R_j = D_j$ / T.
- **Qualification.** Q = { j : $S_j > 0.20$ }. If V = 0 or Q = ∅, return candidate-set rejection.
- **Selection.** Minimum R_j within Q; tie broken by higher S_j; any remaining tie by a predetermined candidate-symmetric neutral procedure.
- **Public output.** Publish E, T, V, T/E, V/E, V/T and, for every candidate, A_j, A_j/T, S_j, D_j, R_j, with Q and the winner or failure state. **A_j/T and V/T are published inseparably from the result; there is no validity threshold on V/T.**
- **Semantics.** The ballot instructions **must** state which reading of the negative mark applies: aversion, or avoidance of viable rivals. The choice determines whether absorbing entry is exploitable (Appendix C) and whether the selection stage collapses to plurality within Q at q ≤ d + 1 (Proposition 8).
- **Non-compensation.** Positive marks determine qualification; negative marks determine the primary ordering inside Q; the two are never netted.
- **Sole qualifier.** If Q = {a}, a wins regardless of R_a.
- **Candidate-set freeze.** Post-commencement withdrawal, death or disqualification handled by a rule fixed in advance.
- **Continuation rule.** Must satisfy: (1) the reversion is not returned to the agenda setter; (2) any exclusion threshold is set well below α; (3) a termination condition is specified.
- **Scope.** Not for contests with fewer than four candidates, multi-winner elections, party lists, referenda or write-ins.

Appendix B. Reporting Template

Field	Meaning	Invalid interpretation
E	Valid submitted ballots	Number of voters who expressed a candidate judgment
T	Ballots with at least one candidate-directed signal	Population turnout
V	Ballots issuing a positive authorization	Number who found some candidate psychologically good
V/T	Share of active participants supplying affirmative direction	Measure of satisfaction with the candidate set; **a validity criterion for the result**
$S_j = A_j$/V	Candidate j's share of issued positive authorizations	Population support or awareness rate
A_j/T	Positive reach among active participants	Preference intensity, or recognition
$R_j = D_j$/T	Explicit rejection incidence among active participants	Share finding j unacceptable, absent independent validation
$1 - R_j$	Share not recording an objection to j	Share approving, accepting, or having evaluated j
Q	Candidates passing the affirmative concentration gate	Complete set of socially acceptable candidates
w	Rule-selected officeholder	Majority, Condorcet, or universal-consent winner

Recommended narrative statement: "Candidate *w* received S_w of issued affirmative authorizations, qualified under the strict 20% rule, and had the lowest recorded rejection among the qualified candidates, with V/T = […]." Any broader mandate claim requires additional evidence.

Appendix C. A Closed-Form Instance of Absorbing Entry

This appendix gives the example referenced in Sections 6.9 and 9.3.2 in a form that can be verified by hand. Every quantity is exact; no simulation is involved.

C.1 Setup

One ideological dimension. Three candidates and, in the second configuration, one entrant:

Candidate	Position
A	−0.6
B	0.0
C	+0.5
X (entrant)	−1.6

The electorate consists of three blocs located at the positions of A, B and C. Every participant authorizes the nearest candidate and issues two rejections, so V = T and $R_j = D_j$ expressed as a share.

Distances from each bloc:

Bloc	to A	to B	to C	to X
L (−0.6)	0.0	0.6	**1.1**	1.0
M (0.0)	0.6	0.0	0.5	**1.6**
R (+0.5)	1.1	0.5	0.0	**2.1**

No two distances from any bloc are equal, so every "two furthest" set is unique and no tie-breaking convention is needed.

C.2 Configuration 1: no entrant

Bloc shares: $\ell = 0.40$, m = 0.24, r = 0.36. Affirmative shares are therefore $S_A = 0.40$, $S_B = 0.24$, $S_C = 0.36$; all exceed 0.20, so Q = {A, B, C}.

Under the aversion reading each bloc rejects its two furthest candidates:

Bloc	Two furthest	Rejects
L	C (1.1), B (0.6)	B, C
M	A (0.6), C (0.5)	A, C
R	A (1.1), B (0.5)	A, B

Rejection rates:

R_A = m + r = 0.24 + 0.36 = **3/5 = 0.60** $R_B = \ell$ + r = 0.40 + 0.36 = **19/25 = 0.76** $R_C = \ell$ + m = 0.40 + 0.24 = **16/25 = 0.64**

The minimum is R_A. **Candidate A wins**, which is also the plurality leader.

C.3 Configuration 2: entrant X, aversion reading

X attracts ε = 0.02 of the electorate from the L bloc — the most peripheral left voters — who now authorize X. Affirmative shares become $S_A = 0.38$, $S_B = 0.24$, $S_C = 0.36$, $S_X = 0.02$. Since 0.02 < 0.20, **X does not qualify**, and Q = {A, B, C} exactly as before.

Rejections now range over four candidates:

Bloc (share)	Two furthest	Rejects	Change
L minus ε (0.38)	C (1.1), X (1.0)	C, X	**B spared**
ε voters (0.02)	authorize X; two furthest of the rest: C (1.1), B (0.6)	B, C	—
M (0.24)	X (1.6), A (0.6)	A, X	**C spared**
R (0.36)	X (2.1), A (1.1)	A, X	**B spared**

Rejection rates:

R_A = 0.24 + 0.36 = **3/5 = 0.60** R_B = 0.02 = **1/50 = 0.02** R_C = 0.38 + 0.02 = **2/5 = 0.40** R_X = 0.38 + 0.24 + 0.36 = **49/50 = 0.98**

The minimum over Q is R_B. **Candidate B wins.**

Candidate B holds the smallest affirmative share of the three qualified candidates (0.24 against 0.38 and 0.36) and did not gain a single authorization from X's entry. The entire change is in the negative channel: X occupies the top rejection slot for the M and R blocs, displacing A from the M bloc's set and B from the R bloc's set, and displacing B from the residual L bloc's set.

C.4 Configuration 3: entrant X, avoidance reading

Under the avoidance reading each voter directs rejections at the qualified rivals she most wants to prevent from winning. X cannot qualify and therefore receives none.

Bloc	Authorizes	Qualified rivals	Rejects (two furthest rivals)
L minus ε	A	B, C	B, C
ε voters	X (unqualified)	A, B, C	C (1.1), B (0.6)
M	B	A, C	A, C
R	C	A, B	A, B

Rejection rates:

R_A = 0.24 + 0.36 = **3/5 = 0.60** R_B = 0.38 + 0.02 + 0.36 = **19/25 = 0.76** R_C = 0.38 + 0.02 + 0.24 = **16/25 = 0.64** R_X = **0**

These are identical to Configuration 1. **Candidate A wins**, and X's entry has no effect on the outcome.

C.5 What the example establishes and what it does not

Establishes. Under the aversion reading, an entrant holding 2% of issued authorizations and incapable of qualifying can change the winner, and can do so in favour of the qualified candidate with the smallest affirmative share. Under the avoidance reading, in the same configuration, the entrant is inert. The difference is produced entirely by the definition written into the ballot instructions.

Does not establish. That the manipulation is profitable in equilibrium, that it is available for any configuration of positions and shares, that the avoidance reading closes it generally, or that entry of this kind is feasible under real ballot-access law. Section 10.6 states the entry game that would settle the first three questions.

Also worth noting. The beneficiary here is the centre candidate, so in this configuration absorbing entry happens to be centripetal. Nothing in the mechanism guarantees that. The displacement operates on whichever candidate occupies a bloc's second rejection slot, and the identity of that candidate depends on the configuration.